

Discounting in LTL

Shaull Almagor¹, Udi Boker², and Orna Kupferman¹

¹ The Hebrew University, Jerusalem, Israel.

² The Interdisciplinary Center, Herzliya, Israel.

Abstract. In recent years, there is growing need and interest in formalizing and reasoning about the quality of software and hardware systems. As opposed to traditional verification, where one handles the question of whether a system satisfies, or not, a given specification, reasoning about quality addresses the question of *how well* the system satisfies the specification. One direction in this effort is to refine the “eventually” operators of temporal logic to *discounting operators*: the satisfaction value of a specification is a value in $[0, 1]$, where the longer it takes to fulfill eventuality requirements, the smaller the satisfaction value is.

In this paper we introduce an augmentation by discounting of Linear Temporal Logic (LTL), and study it, as well as its combination with propositional quality operators. We show that one can augment LTL with an arbitrary set of discounting functions, while preserving the decidability of the model-checking problem. Further augmenting the logic with unary propositional quality operators preserves decidability, whereas adding an average-operator makes some problems undecidable. We also discuss the complexity of the problem, as well as various extensions.

1 Introduction

One of the main obstacles to the development of complex hardware and software systems lies in ensuring their correctness. A successful paradigm addressing this obstacle is *temporal-logic model checking* – given a mathematical model of the system and a temporal-logic formula that specifies a desired behavior of it, decide whether the model satisfies the formula [5]. Correctness is Boolean: a system can either satisfy its specification or not satisfy it. The richness of today’s systems, however, justifies specification formalisms that are *multi-valued*. The multi-valued setting arises directly in systems with quantitative aspects (multi-valued / probabilistic / fuzzy) [9,10,11,16,23], but is applied also with respect to Boolean systems, where it originates from the semantics of the specification formalism itself [1,7].

When considering the *quality* of a system, satisfying a specification should no longer be a yes/no matter. Different ways of satisfying a specification should induce different levels of quality, which should be reflected in the output of the verification procedure. Consider for example the specification $G(\text{request} \rightarrow F(\text{response_grant} \vee \text{response_deny}))$ (“every request is eventually responded, with either a grant or a denial”). There should be a difference between a computation that satisfies it with responses generated soon after requests and one that satisfies it with long waits. Moreover, there may be a difference between grant and deny responses, or cases in which no request is issued. The issue of generating high-quality hardware and software systems attracts a lot of attention [13,26]. Quality, however, is traditionally viewed as an art, or as an amorphous ideal. In [1], we introduced an approach for formalizing quality. Using it, a user can specify quality formally, according to the importance he gives to components such as security, maintainability, runtime, and more, and then can formally reason about the quality of software.

As the example above demonstrates, we can distinguish between two aspects of the quality of satisfaction. The first, to which we refer as “temporal quality” concerns the waiting time to satisfaction of eventualities. The second, to which we refer as “propositional quality” concerns prioritizing related components of the specification. Propositional quality was studied in [1]. In this paper we study temporal quality as well as the combinations of both aspects. One may try to reduce temporal quality to propositional quality by a repeated use of the X (“next”) operator or by a use of bounded (prompt) eventualities [2,3]. Both approaches, however, partition the future into finitely many zones and are limited: correctness of LTL is Boolean, and thus has inherent dichotomy between satisfaction and dissatisfaction. On the other hand, the distinction between “near” and “far” is not dichotomous. This suggests that in order to formalize temporal quality, one must extend LTL to an unbounded setting. Realizing this, researchers have suggested to augment temporal logics with *future discounting* [8]. In the discounted setting, the satisfaction value of specifications is a numerical value, and it depends, according to some discounting function, on the time waited for eventualities to get satisfied.

In this paper we add discounting to Linear Temporal Logic (LTL), and study it, as well as its combination with propositional quality operators. We introduce $\text{LTL}^{\text{disc}}[\mathcal{D}]$ – an augmentation by discounting of LTL. The logic

$\text{LTL}^{\text{disc}}[\mathcal{D}]$ is actually a family of logics, each parameterized by a set \mathcal{D} of discounting functions – strictly decreasing functions from \mathbb{N} to $[0, 1]$ that tend to 0 (e.g., linear decaying, exponential decaying, etc.). $\text{LTL}^{\text{disc}}[\mathcal{D}]$ includes a discounting-“until” (U_η) operator, parameterized by a function $\eta \in \mathcal{D}$. We solve the model-checking threshold problem for $\text{LTL}^{\text{disc}}[\mathcal{D}]$: given a Kripke structure \mathcal{K} , an $\text{LTL}^{\text{disc}}[\mathcal{D}]$ formula φ and a threshold $t \in [0, 1]$, the algorithm decides whether the satisfaction value of φ in \mathcal{K} is at least t .

In the Boolean setting, the automata-theoretic approach has proven to be very useful in reasoning about LTL specifications. The approach is based on translating LTL formulas to nondeterministic Büchi automata on infinite words [28]. Applying this approach to the discounted setting, which gives rise to infinitely many satisfaction values, poses a big algorithmic challenge: model-checking algorithms, and in particular those that follow the automata-theoretic approach, are based on an exhaustive search, which cannot be simply applied when the domain becomes infinite. A natural relevant extension to the automata-theoretic approach is to translate formulas to *weighted automata* [22]. Unfortunately, these extensively-studied models are complicated and many problems become undecidable for them [15]. We show that for threshold problems, we can translate $\text{LTL}^{\text{disc}}[\mathcal{D}]$ formulas into (Boolean) nondeterministic Büchi automata, with the property that the automaton accepts a lasso computation iff the formula attains a value above the threshold on that computation. Our algorithm relies on the fact that the language of an automaton is non-empty iff there is a lasso witness for the non-emptiness. We cope with the infinitely many possible satisfaction values by using the discounting behavior of the eventualities and the given threshold in order to partition the state space into a finite number of classes. The complexity of our algorithm depends on the discounting functions used in the formula. We show that for standard discounting functions, such as exponential decaying, the problem is PSPACE-complete – not more complex than standard LTL. The fact our algorithm uses Boolean automata also enables us to suggest a solution for threshold satisfiability, and to give a partial solution to threshold synthesis. In addition, it allows to adapt the heuristics and tools that exist for Boolean automata.

Before we continue to describe our contribution, let us review existing work on discounting. The notion of discounting has been studied in several fields, such as economy, game-theory, and Markov decision processes [25]. In the area of formal verification, it was suggested in [8] to augment the μ -calculus with discounting operators. The discounting suggested there is exponential; that is, with each iteration, the satisfaction value of the formula decreases by a multiplicative factor in $(0, 1]$. Algorithmically, [8] shows how to evaluate discounted μ -calculus formulas with arbitrary precision. Formulas of LTL can be translated to the μ -calculus, thus [8] can be used in order to approximately model-check discounted-LTL formulas. However, the translation from LTL to the μ -calculus involves an exponential blowup [6] (and is complicated), making this approach inefficient. Moreover, our approach allows for arbitrary discounting functions, and the algorithm returns an exact solution to the threshold model-checking problem, which is more difficult than the approximation problem.

Closer to our work is [7], where CTL is augmented with discounting and weighted-average operators. The motivation in [7] is to introduce a logic whose semantics is not too sensitive to small perturbations in the model. Accordingly, formulas are evaluated on weighted-systems or on Markov-chains. Adding discounting and weighted-average operators to CTL preserves its appealing complexity, and the model-checking problem for the augmented logic can be solved in polynomial time. As is the case in the traditional, Boolean, semantics, the expressive power of discounted CTL is limited. The fact the same combination, of discounting and weighted-average operators, leads to undecidability in the context of LTL witnesses the technical challenges of the $\text{LTL}^{\text{disc}}[\mathcal{D}]$ setting.

Perhaps closest to our approach is [19], where a version of discounted-LTL was introduced. Semantically, there are two main differences between the logics. The first is that [19] uses discounted sum, while we interpret discounting without accumulation, and the second is that the discounting there replaces the standard temporal operators, so all eventualities are discounted. As discounting functions tend to 0, this strictly restricts the expressive power of the logic, and one cannot specify traditional eventualities in it. On the positive side, it enables a clean algebraic characterization of the semantics, and indeed the contribution in [19] is a comprehensive study of the mathematical properties of the logic. Yet, [19] does not study algorithmic questions about the logic. We, on the other hand, focus on the algorithmic properties of the logic, and specifically on the model-checking problem.

Let us now return to our contribution. After introducing $\text{LTL}^{\text{disc}}[\mathcal{D}]$ and studying its model-checking problem, we augment $\text{LTL}^{\text{disc}}[\mathcal{D}]$ with propositional quality operators. Beyond the operators \min , \max , and \neg , which are already present, two basic propositional quality operators are the multiplication of an $\text{LTL}^{\text{disc}}[\mathcal{D}]$ formula by a constant in $[0, 1]$, and the averaging between the satisfaction values of two $\text{LTL}^{\text{disc}}[\mathcal{D}]$ formulas [1]. We show that while the first

extension does not increase the expressive power of $LTL^{\text{disc}}[\mathcal{D}]$ or its complexity, the latter causes some problems (e.g. validity) to become undecidable. In fact, things become undecidable even if we allow averaging in combination with a single discounting function. Recall that this is in contrast with the extension of discounted CTL with an average operator, where the complexity of the model-checking problem stays polynomial [7].

We consider additional extensions of $LTL^{\text{disc}}[\mathcal{D}]$. First, we study a variant of the discounting-eventually operators in which we allow the discounting to tend to arbitrary values in $[0, 1]$ (rather than to 0). This captures the intuition that we are not always pessimistic about the future, but can be, for example, ambivalent about it, by tending to $\frac{1}{2}$. We show that all our results hold under this extension. Second, we add to $LTL^{\text{disc}}[\mathcal{D}]$ *past* operators and their discounting versions (specifically, we allow a discounting-“since” operator, and its dual). In the traditional semantics, past operators enable clean specifications of many interesting properties, make the logic exponentially more succinct, and can still be handled within the same complexity bounds [17,18]. We show that the same holds for the discounted setting. Finally, we show how $LTL^{\text{disc}}[\mathcal{D}]$ and algorithms for it can be used also for reasoning about weighted systems.

Due to lack of space, the full proofs appear in the appendix.

2 The Logic $LTL^{\text{disc}}[\mathcal{D}]$

The linear temporal logic $LTL^{\text{disc}}[\mathcal{D}]$ generalizes LTL by adding discounting temporal operators. The logic is actually a family of logics, each parameterized by a set \mathcal{D} of discounting functions.

Let $\mathbb{N} = \{0, 1, \dots\}$. A function $\eta : \mathbb{N} \rightarrow [0, 1]$ is a *discounting function* if $\lim_{i \rightarrow \infty} \eta(i) = 0$, and η is strictly monotonic-decreasing. Examples for natural discounting functions are $\eta(i) = \lambda^i$, for some $\lambda \in (0, 1)$, and $\eta(i) = \frac{1}{i+1}$.

Given a set of discounting functions \mathcal{D} , we define the logic $LTL^{\text{disc}}[\mathcal{D}]$ as follows. The syntax of $LTL^{\text{disc}}[\mathcal{D}]$ adds to LTL the operator $\varphi U_\eta \psi$ (discounting-Until), for every function $\eta \in \mathcal{D}$. Thus, the syntax is given by the following grammar, where p ranges over the set AP of atomic propositions and $\eta \in \mathcal{D}$.

$$\varphi := \text{True} \mid p \mid \neg\varphi \mid \varphi \vee \varphi \mid X\varphi \mid \varphi U \varphi \mid \varphi U_\eta \varphi.$$

The semantics of $LTL^{\text{disc}}[\mathcal{D}]$ is defined with respect to a *computation* $\pi = \pi^0, \pi^1, \dots \in (2^{AP})^\omega$. Given a computation π and an $LTL^{\text{disc}}[\mathcal{D}]$ formula φ , the truth value of φ in π is a value in $[0, 1]$, denoted $\llbracket \pi, \varphi \rrbracket$. The value is defined by induction on the structure of φ as follows, where $\pi^i = \pi_i, \pi_{i+1}, \dots$.

$$\begin{aligned} - \llbracket \pi, \text{True} \rrbracket &= 1. & - \llbracket \pi, \varphi \vee \psi \rrbracket &= \max\{\llbracket \pi, \varphi \rrbracket, \llbracket \pi, \psi \rrbracket\}. \\ - \llbracket \pi, p \rrbracket &= \begin{cases} 1 & \text{if } p \in \pi_0, \\ 0 & \text{if } p \notin \pi_0. \end{cases} & - \llbracket \pi, \neg\varphi \rrbracket &= 1 - \llbracket \pi, \varphi \rrbracket. \\ - \llbracket \pi, X\varphi \rrbracket &= \llbracket \pi^1, \varphi \rrbracket. \\ - \llbracket \pi, \varphi U \psi \rrbracket &= \sup_{i \geq 0} \{\min\{\llbracket \pi^i, \psi \rrbracket, \min_{0 \leq j < i} \{\llbracket \pi^j, \varphi \rrbracket\}\}\}. \\ - \llbracket \pi, \varphi U_\eta \psi \rrbracket &= \sup_{i \geq 0} \{\min\{\eta(i) \llbracket \pi^i, \psi \rrbracket, \min_{0 \leq j < i} \{\eta(j) \llbracket \pi^j, \varphi \rrbracket\}\}\}. \end{aligned}$$

The intuition is that events that happen in the future have a lower influence, and the rate by which this influence decreases depends on the function η .³ For example, the satisfaction value of a formula $\varphi U_\eta \psi$ in a computation π depends on the best (supremum) value that ψ can get along the entire computation, while considering the discounted satisfaction of ψ at a position i , as a result of multiplying it by $\eta(i)$, and the same for the value of φ in the prefix leading to the i -th position.

We add the standard abbreviations $F\varphi \equiv \text{True} U \varphi$, and $G\varphi \equiv \neg F \neg \varphi$, as well as their quantitative counterparts: $F_\eta \varphi \equiv \text{True} U_\eta \varphi$, and $G_\eta \varphi \equiv \neg F_\eta \neg \varphi$. We denote by $|\varphi|$ the number of subformulas of φ .

A computation of the form $\pi = u \cdot v^\omega$, for $u, v \in (2^{AP})^*$, with $v \neq \epsilon$, is called a *lasso computation*. We observe that since a specific lasso computation has only finitely many distinct suffixes, the \inf and \sup in the semantics of $LTL^{\text{disc}}[\mathcal{D}]$ can be replaced with \min and \max , respectively, when applied to lasso computations.

The semantics is extended to *Kripke structures* by taking the path that admits the lowest satisfaction value. Formally, for a Kripke structure \mathcal{K} and an $LTL^{\text{disc}}[\mathcal{D}]$ formula φ we have that $\llbracket \mathcal{K}, \varphi \rrbracket = \inf \{\llbracket \pi, \varphi \rrbracket : \pi \text{ is a computation of } \mathcal{K}\}$.

³ Observe that in our semantics the satisfaction value of future events tends to 0. One may think of scenarios where future events are discounted towards another value in $[0, 1]$ (e.g. discounting towards $\frac{1}{2}$ as ambivalence regarding the future). We address this in Section 5.3.

Example 1. Consider a lossy-disk: every moment in time there is a chance that some bit would flip its value. Fixing flips is done by a global error-correcting procedure. This procedure manipulates the entire content of the disk, such that initially it causes more errors in the disk, but the longer it runs, the more bits it fixes.

Let *init* and *terminate* be atomic propositions indicating when the error-correcting procedure is initiated and terminated, respectively. The quality of the disk (that is, a measure of the amount of correct bits) can be specified by the formula $\varphi = \text{GF}_\eta(\text{init} \wedge \neg \text{F}_\mu \text{terminate})$ for some appropriate discounting functions η and μ . Intuitively, φ gets a higher satisfaction value the shorter the waiting time is between initiations of the error-correcting procedure, and the longer the procedure runs (that is, not terminated) in between these initiations. Note that the “worst case” nature of $\text{LTL}^{\text{disc}}[\mathcal{D}]$ fits here. For instance, running the procedure for a very short time, even once, will cause many errors.

3 $\text{LTL}^{\text{disc}}[\mathcal{D}]$ Model Checking

In the Boolean setting, the model-checking problem asks, given an LTL formula φ and a Kripke structure \mathcal{K} , whether $\llbracket \mathcal{K}, \varphi \rrbracket = \text{True}$. In the quantitative setting, the model-checking problem is to compute $\llbracket \mathcal{K}, \varphi \rrbracket$, where φ is now an $\text{LTL}^{\text{disc}}[\mathcal{D}]$ formula. A simpler version of this problem is the threshold model-checking problem: given φ , \mathcal{K} , and a threshold $v \in [0, 1]$, decide whether $\llbracket \mathcal{K}, \varphi \rrbracket \geq v$. In this section we show how we can solve the latter.

Our solution uses the automata-theoretic approach, and consists of the following steps. We start by translating φ and v to an alternating weak automaton $\mathcal{A}_{\varphi, v}$ such that $L(\mathcal{A}_{\varphi, v}) \neq \emptyset$ iff there exists a computation π such that $\llbracket \pi, \varphi \rrbracket > v$. The challenge here is that φ has infinitely many satisfaction values, naively implying an infinite-state automaton. We show that using the threshold and the discounting behavior of the eventualities, we can restrict attention to a finite resolution of satisfaction values, enabling the construction of a finite automaton. Complexity-wise, the size of $\mathcal{A}_{\varphi, v}$ depends on the functions in \mathcal{D} . In Section 3.3, we analyze the complexity for the case of exponential-discounting functions.

The second step is to construct a nondeterministic Büchi automaton \mathcal{B} that is equivalent to $\mathcal{A}_{\varphi, v}$. In general, alternation removal might involve an exponential blowup in the state space [21]. We show, by a careful analysis of $\mathcal{A}_{\varphi, v}$, that we can remove its alternation while only having a polynomial state blowup.

We complete the model-checking procedure by composing the nondeterministic Büchi automaton \mathcal{B} with the Kripke structure \mathcal{K} , as done in the traditional, automata-based, model-checking procedure.

The complexity of model-checking an $\text{LTL}^{\text{disc}}[\mathcal{D}]$ formula depends on the discounting functions in \mathcal{D} . Intuitively, the faster the discounting tends to 0, the less states there will be. For exponential-discounting, we show that the complexity is NLOGSPACE in the system (the Kripke structure) and PSPACE in the specification (the $\text{LTL}^{\text{disc}}[\mathcal{D}]$ formula and the threshold), staying in the same complexity classes of standard LTL model-checking.

We conclude the section by showing how to use the generated nondeterministic Büchi automaton for addressing threshold satisfiability and synthesis.

3.1 Alternating Weak Automata

For a given set X , let $\mathcal{B}^+(X)$ be the set of positive Boolean formulas over X (i.e., Boolean formulas built from elements in X using \wedge and \vee), where we also allow the formulas **True** and **False**. For $Y \subseteq X$, we say that Y *satisfies* a formula $\theta \in \mathcal{B}^+(X)$ iff the truth assignment that assigns *true* to the members of Y and assigns *false* to the members of $X \setminus Y$ satisfies θ . An *alternating Büchi automaton on infinite words* is a tuple $\mathcal{A} = \langle \Sigma, Q, q_{\text{in}}, \delta, \alpha \rangle$, where Σ is the input alphabet, Q is a finite set of states, $q_{\text{in}} \in Q$ is an initial state, $\delta : Q \times \Sigma \rightarrow \mathcal{B}^+(Q)$ is a transition function, and $\alpha \subseteq Q$ is a set of accepting states. We define runs of \mathcal{A} by means of (possibly) infinite DAGs (directed acyclic graphs). A run of \mathcal{A} on a word $w = \sigma_0 \cdot \sigma_1 \cdot \dots \in \Sigma^\omega$ is a (possibly) infinite DAG $\mathcal{G} = \langle V, E \rangle$ satisfying the following (note that there may be several runs of \mathcal{A} on w).

- $V \subseteq Q \times \mathbb{N}$ is as follows. Let $Q_l \subseteq Q$ denote all states in level l . Thus, $Q_l = \{q : \langle q, l \rangle \in V\}$. Then, $Q_0 = \{q_{\text{in}}\}$, and Q_{l+1} satisfies $\bigwedge_{q \in Q_l} \delta(q, \sigma_l)$.
- For every $l \in \mathbb{N}$, Q_l is minimal with respect to containment.
- $E \subseteq \bigcup_{l \geq 0} (Q_l \times \{l\}) \times (Q_{l+1} \times \{l+1\})$ is such that for every state $q \in Q_l$, the set $\{q' \in Q_{l+1} : E(\langle q, l \rangle, \langle q', l+1 \rangle)\}$ satisfies $\delta(q, \sigma_l)$.

Thus, the root of the DAG contains the initial state of the automaton, and the states associated with nodes in level $l + 1$ satisfy the transitions from states corresponding to nodes in level l . The run \mathcal{G} accepts the word w if all its infinite paths satisfy the acceptance condition α . Thus, in the case of Büchi automata, all the infinite paths have infinitely many nodes $\langle q, l \rangle$ such that $q \in \alpha$ (it is not hard to prove that every infinite path in \mathcal{G} is part of an infinite path starting in level 0). A word w is accepted by \mathcal{A} if there is a run that accepts it. The language of \mathcal{A} , denoted $L(\mathcal{A})$, is the set of infinite words that \mathcal{A} accepts.

When the formulas in the transition function of \mathcal{A} contain only disjunctions, then \mathcal{A} is nondeterministic, and its runs are DAGs of width 1, where at each level there is a single node.

The alternating automaton \mathcal{A} is *weak*, denoted AWA, if its state space Q can be partitioned into sets Q_1, \dots, Q_k , such that the following hold: First, for every $1 \leq i \leq k$ either $Q_i \subseteq \alpha$, in which case we say that Q_i is an accepting set, or $Q_i \cap \alpha = \emptyset$, in which case we say that Q_i is rejecting. Second, there is a partial-order \leq over the sets, and for every $1 \leq i, j \leq k$, if $q \in Q_i$, $s \in Q_j$, and $s \in \delta(q, \sigma)$ for some $\sigma \in \Sigma$, then $Q_j \leq Q_i$. Thus, transitions can lead only to states that are smaller in the partial order. Consequently, each run of an AWA eventually gets trapped in a set Q_i and is accepting iff this set is accepting.

3.2 From $\text{LTL}^{\text{disc}}[\mathcal{D}]$ to AWA

Our model-checking algorithm is based on translating an $\text{LTL}^{\text{disc}}[\mathcal{D}]$ formula φ to an AWA. Intuitively, the states of the AWA correspond to assertions of the form $\psi > t$ or $\psi < t$ for every subformula ψ of φ , and for certain thresholds $t \in [0, 1]$. A lasso computation is then accepted from state $\psi > t$ iff $\llbracket \pi, \psi \rrbracket > t$. The assumption about the computation being a lasso is needed only for the “only if” direction, and it does not influence the proof’s generality since the language of an automaton is non-empty iff there is a lasso witness for its non-emptiness. By setting the initial state to $\varphi > v$, we are done.

Defining the appropriate transition function for the AWA follows the semantics of $\text{LTL}^{\text{disc}}[\mathcal{D}]$ in the expected manner. A naive construction, however, yields an infinite-state automaton (even if we only expand the state space on-the-fly, as discounting formulas can take infinitely many satisfaction values). As can be seen in the proof of Theorem 1, the “problematic” transitions are those that involve the discounting operators. The key observation is that, given a threshold v and a computation π , when evaluating a discounted operator on π , one can restrict attention to two cases: either the satisfaction value of the formula goes below v , in which case this happens after a bounded prefix, or the satisfaction value always remains above v , in which case we can replace the discounted operator with a Boolean one. This observation allows us to expand only a finite number of states on-the-fly.

Before describing the construction of the AWA, we need the following lemma, which reduces an extreme satisfaction of an $\text{LTL}^{\text{disc}}[\mathcal{D}]$ formula, meaning satisfaction with a value of either 0 or 1, to a Boolean satisfaction of an LTL formula. The proof proceeds by induction on the structure of the formulas.

Lemma 1. *Given an $\text{LTL}^{\text{disc}}[\mathcal{D}]$ formula φ , there exist LTL formulas φ^+ and $\varphi^{<1}$ such that $|\varphi^+|$ and $|\varphi^{<1}|$ are both $O(|\varphi|)$ and the following hold for every computation π .*

1. *If $\llbracket \pi, \varphi \rrbracket > 0$ then $\pi \models \varphi^+$, and if $\llbracket \pi, \varphi \rrbracket < 1$ then $\pi \models \varphi^{<1}$.*
2. *If π is a lasso, then if $\pi \models \varphi^+$ then $\llbracket \pi, \varphi \rrbracket > 0$ and if $\pi \models \varphi^{<1}$ then $\llbracket \pi, \varphi \rrbracket < 1$.*

Henceforth, given an $\text{LTL}^{\text{disc}}[\mathcal{D}]$ formula φ , we refer to φ^+ as in Lemma 1.

Consider an $\text{LTL}^{\text{disc}}[\mathcal{D}]$ formula φ . By Lemma 1, if there exists a computation π such that $\llbracket \pi, \varphi \rrbracket > 0$, then φ^+ is satisfiable. Conversely, since φ^+ is a Boolean LTL formula, then by [27] we know that φ^+ is satisfiable iff there exists a lasso computation π that satisfies it, in which case $\llbracket \pi, \varphi \rrbracket > 0$. We conclude with the following.

Corollary 1. *Consider an $\text{LTL}^{\text{disc}}[\mathcal{D}]$ formula φ . There exists a computation π such that $\llbracket \pi, \varphi \rrbracket > 0$ iff there exists a lasso computation π' such that $\llbracket \pi', \varphi \rrbracket > 0$, in which case $\pi' \models \varphi^+$ as well.*

Remark 1. The curious reader may wonder why we do not prove that $\llbracket \pi, \varphi \rrbracket > 0$ iff $\pi \models \varphi^+$ for every computation π . As it turns out, a translation that is valid also for computations with no period is not always possible. For example, as is the case with the prompt-eventuality operator of [14], the formula $\varphi = \text{G}(\text{F}_{\eta}p)$ is such that the set of computations π with $\llbracket \pi, \varphi \rrbracket > 0$ is not ω -regular, thus one cannot hope to define an LTL formula φ^+ .

We start with some definitions. For a function $f : \mathbb{N} \rightarrow [0, 1]$ and for $k \in \mathbb{N}$, we define $f^{+k} : \mathbb{N} \rightarrow [0, 1]$ as follows. For every $i \in \mathbb{N}$ we have that $f^{+k}(i) = f(i + k)$.

Let φ be an $\text{LTL}^{\text{disc}}[\mathcal{D}]$ formula over AP . We define the *extended closure* of φ , denoted $xcl(\varphi)$, to be the set of all the formulas ψ of the following *classes*:

1. ψ is a subformula of φ .
2. ψ is a subformula of θ^+ or $\neg\theta^+$, where θ is a subformula of φ .
3. ψ is of the form $\theta_1 \text{U}_{\eta+k} \theta_2$ for $k \in \mathbb{N}$, where $\theta_1 \text{U}_{\eta} \theta_2$ is a subformula of φ .

Observe that $xcl(\varphi)$ may be infinite, and that it has both $\text{LTL}^{\text{disc}}[\mathcal{D}]$ formulas (from Classes 1 and 3) and LTL formulas (from Class 2).

Theorem 1. *Given an $\text{LTL}^{\text{disc}}[\mathcal{D}]$ formula φ and a threshold $v \in [0, 1]$, there exists an AWA $\mathcal{A}_{\varphi,v}$ such that for every computation π the following hold.*

1. If $\llbracket \pi, \varphi \rrbracket > v$, then $\mathcal{A}_{\varphi,v}$ accepts π .
2. If $\mathcal{A}_{\varphi,v}$ accepts π and π is a lasso computation, then $\llbracket \pi, \varphi \rrbracket > v$.

Proof: We construct $\mathcal{A}_{\varphi,v} = \langle Q, 2^{AP}, Q_0, \delta, \alpha \rangle$ as follows.

The state space Q consists of two types of states. Type-1 states are assertions of the form $(\psi > t)$ or $(\psi < t)$, where $\psi \in xcl(\varphi)$ is of Class 1 or 3 and $t \in [0, 1]$. Type-2 states correspond to LTL formulas of Class 2. Let S be the set of Type-1 and Type-2 states for all $\psi \in xcl(\varphi)$ and thresholds $t \in [0, 1]$. Then, Q is the subset of S constructed on-the-fly according to the transition function defined below. We later show that Q is indeed finite.

The transition function $\delta : Q \times 2^{AP} \rightarrow \mathcal{B}^+(Q)$ is defined as follows. For Type-2 states, the transitions are as in the standard translation from LTL to AWA [27] (see Appendix A.2 for details). For the other states, we define the transitions as follows. Let $\sigma \in 2^{AP}$.

$$\begin{aligned}
- \delta((\text{True} > t), \sigma) &= \begin{cases} \text{True} & \text{if } t < 1, \\ \text{False} & \text{if } t = 1. \end{cases} \\
- \delta((\text{False} > t), \sigma) &= \text{False}. \\
- \delta((\text{True} < t), \sigma) &= \text{False}. \\
- \delta((\text{False} < t), \sigma) &= \begin{cases} \text{True} & \text{if } t > 0, \\ \text{False} & \text{if } t = 0. \end{cases} \\
- \delta((p > t), \sigma) &= \begin{cases} \text{True} & \text{if } p \in \sigma \text{ and } t < 1, \\ \text{False} & \text{otherwise.} \end{cases} \\
- \delta((p < t), \sigma) &= \begin{cases} \text{False} & \text{if } p \in \sigma \text{ or } t = 0, \\ \text{True} & \text{otherwise.} \end{cases} \\
- \delta((\psi_1 \vee \psi_2 > t), \sigma) &= \delta((\psi_1 > t), \sigma) \vee \delta((\psi_2 > t), \sigma). \\
- \delta((\psi_1 \vee \psi_2 < t), \sigma) &= \delta((\psi_1 < t), \sigma) \wedge \delta((\psi_2 < t), \sigma). \\
- \delta((\neg\psi_1 > t), \sigma) &= \delta((\psi_1 < 1 - t), \sigma) \\
- \delta((\neg\psi_1 < t), \sigma) &= \delta((\psi_1 > 1 - t), \sigma). \\
- \delta((X\psi_1 > t), \sigma) &= (\psi_1 > t). \\
- \delta((X\psi_1 < t), \sigma) &= (\psi_1 < t). \\
- \delta((\psi_1 \text{U}\psi_2 > t), \sigma) &= \begin{cases} \delta((\psi_2 > t), \sigma) \vee [\delta((\psi_1 > t), \sigma) \wedge (\psi_1 \text{U}\psi_2 > t)] & \text{if } 0 < t < 1, \\ \text{False} & \text{if } t \geq 1, \\ \delta(((\psi_1 \text{U}\psi_2)^+), \sigma) & \text{if } t = 0. \end{cases} \\
- \delta((\psi_1 \text{U}\psi_2 < t), \sigma) &= \begin{cases} \delta((\psi_2 < t), \sigma) \wedge [\delta((\psi_1 < t), \sigma) \vee (\psi_1 \text{U}\psi_2 < t)] & \text{if } 0 < t \leq 1, \\ \text{True} & \text{if } t > 1, \\ \text{False} & \text{if } t = 0. \end{cases} \\
- \delta((\psi_1 \text{U}_{\eta}\psi_2 > t), \sigma) &= \begin{cases} \delta((\psi_2 > \frac{t}{\eta(0)}), \sigma) \vee [\delta((\psi_1 > \frac{t}{\eta(0)}), \sigma) \wedge (\psi_1 \text{U}_{\eta+1}\psi_2 > t)] & \text{if } 0 < \frac{t}{\eta(0)} < 1, \\ \text{False} & \text{if } \frac{t}{\eta(0)} \geq 1, \\ \delta(((\psi_1 \text{U}_{\eta}\psi_2)^+), \sigma) & \text{if } \frac{t}{\eta(0)} = 0 \text{ (i.e., } t = 0). \end{cases}
\end{aligned}$$

$$- \delta((\psi_1 U_\eta \psi_2 < t), \sigma) = \begin{cases} \delta((\psi_2 < \frac{t}{\eta(0)}), \sigma) \wedge [\delta((\psi_1 < \frac{t}{\eta(0)}), \sigma) \vee (\psi_1 U_{\eta+1} \psi_2 < t)] & \text{if } 0 < \frac{t}{\eta(0)} \leq 1, \\ \text{True} & \text{if } \frac{t}{\eta(0)} > 1, \\ \text{False} & \text{if } \frac{t}{\eta(0)} = 0 \text{ (i.e., } t = 0). \end{cases}$$

We provide some intuition for the more complex parts of the transition function: consider, for example, the transition $\delta((\psi_1 U_\eta \psi_2 > t), \sigma)$. Since η is decreasing, the highest possible satisfaction value for $\psi_1 U_\eta \psi_2$ is $\eta(0)$. Thus, if $\eta(0) \leq t$ (equivalently, $\frac{t}{\eta(0)} \geq 1$), then it cannot hold that $\psi_1 U_\eta \psi_2 > t$, so the transition is to **False**. If $t = 0$, then we only need to ensure that the satisfaction value of $\psi_1 U_\eta \psi_2$ is not 0. To do so, we require that $(\psi_1 U_\eta \psi_2)^+$ is satisfied. By Corollary 1, this is equivalent to the satisfiability of the former. So the transition is identical to that of the state $(\psi_1 U_\eta \psi_2)^+$. Finally, if $0 < t < \eta(0)$, then (slightly abusing notation) the assertion $\psi_1 U_\eta \psi_2 > t$ is true if either $\eta(0)\psi_2 > t$ is true, or both $\eta(0)\psi_1 > t$ and $\psi_1 U_{\eta+1} \psi_2 > t$ are true.

The initial state of $\mathcal{A}_{\varphi, v}$ is $(\varphi > v)$. The accepting states are these of the form $(\psi_1 U \psi_2 < t)$, as well as accepting states that arise in the standard translation of Boolean LTL to AWA (in Type-2 states). Note that each path in the run of $\mathcal{A}_{\varphi, v}$ eventually gets trapped in a single state. Thus, $\mathcal{A}_{\varphi, v}$ is indeed an AWA. The intuition behind the acceptance condition is as follows. Getting trapped in state of the form $(\psi_1 U \psi_2 < t)$ is allowed, as the eventuality is satisfied with value 0. On the other hand, getting stuck in other states (or Type-1) is not allowed, as they involve eventualities that are not satisfied in the threshold promised for them.

This concludes the definition of $\mathcal{A}_{\varphi, v}$. Finally, observe that while the construction as described above is infinite (indeed, uncountable), only finitely many states are reachable from the initial state $(\varphi > v)$, and we can compute these states in advance. Intuitively, it follows from the fact that once the proportion between t and $\eta(i)$ goes above 1, for Type-1 states associated with threshold t and sub formulas with a discounting function η , we do not have to generate new states.

A detailed proof of \mathcal{A} 's finiteness and correctness is given in the appendix. \square

Since $\mathcal{A}_{\varphi, v}$ is a Boolean automaton, then its language is not empty iff it accepts a lasso computation. Combining this observation with Theorem 1, we conclude with the following.

Corollary 2. *For an $\text{LTL}^{\text{disc}}[\mathcal{D}]$ formula φ and a threshold $v \in [0, 1]$, it holds that $L(\mathcal{A}_{\varphi, v}) \neq \emptyset$ iff there exists a computation π such that $\llbracket \pi, \varphi \rrbracket > v$.*

3.3 Exponential Discounting

The size of the AWA generated as per Theorem 1 depends on the discounting functions. In this section, we analyze its size for the class of *exponential discounting* functions, showing that it is singly exponential in the specification formula and in the threshold. This class is perhaps the most common class of discounting functions, as it describes what happens in many natural processes (e.g., temperature change, capacitor charge, effective interest rate, etc.) [8,25].

For $\lambda \in (0, 1)$ we define the *exponential-discounting* function $\exp_\lambda : \mathbb{N} \rightarrow [0, 1]$ by $\exp_\lambda(i) = \lambda^i$. For the purpose of this section, we restrict to $\lambda \in (0, 1) \cap \mathbb{Q}$. Let $E = \{\exp_\lambda : \lambda \in (0, 1) \cap \mathbb{Q}\}$, and consider the logic $\text{LTL}^{\text{disc}}[E]$.

For an $\text{LTL}^{\text{disc}}[E]$ formula φ we define the set $F(\varphi)$ to be $\{\lambda_1, \dots, \lambda_k : \text{the operator } U_{\exp_\lambda} \text{ appears in } \varphi\}$. Let $|\langle \varphi \rangle|$ be the length of the description of φ . That is, in addition to $|\varphi|$, we include in $|\langle \varphi \rangle|$ the length, in bits, of describing $F(\varphi)$.

Theorem 2. *Given an $\text{LTL}^{\text{disc}}[E]$ formula φ and a threshold $v \in [0, 1] \cap \mathbb{Q}$, there exists an AWA $\mathcal{A}_{\varphi, v}$ such that for every computation π the following hold.*

1. *If $\llbracket \pi, \varphi \rrbracket > v$, then $\mathcal{A}_{\varphi, v}$ accepts π .*
2. *If $\mathcal{A}_{\varphi, v}$ accepts π and π is a lasso computation, then $\llbracket \pi, \varphi \rrbracket > v$.*

Furthermore, the number of states of $\mathcal{A}_{\varphi, v}$ is singly exponential in $|\langle \varphi \rangle|$ and in the description of v .

The proof follows from the following observation. Let $\lambda \in (0, 1)$ and $v \in (0, 1)$. When discounting by \exp_λ , the number of states in the AWA constructed as per Theorem 1 is proportional to the maximal number i such that $\lambda^i > v$, which is at most $\log_\lambda v = \frac{\log v}{\log \lambda}$, which is polynomial in the description length of v and λ . A similar (yet more complicated) consideration is applied for the setting of multiple discounting functions and negations.

3.4 From $\mathcal{A}_{\varphi,v}$ to an NBA

Every AWA can be translated to an equivalent nondeterministic Büchi automaton (NBA, for short), yet the state blowup might be exponential BKR10,MH84. By carefully analyzing the AWA $\mathcal{A}_{\varphi,v}$ generated in Theorem 1, we show that it can be translated to an NBA with only a polynomial blowup.

The idea behind our complexity analysis is as follows. Translating an AWA to an NBA involves alternation removal, which proceeds by keeping track of entire levels in a run-DAG. Thus, a run of the NBA corresponds to a sequence of subsets of Q . The key to the reduced state space is that the number of such subsets is only $|Q|^{O(|\varphi|)}$ and not $2^{|Q|}$. To see why, consider a subset S of the states of \mathcal{A} . We say that S is *minimal* if it does not include two states of the form $\varphi < t_1$ and $\varphi < t_2$, for $t_1 < t_2$, nor two states of the form $\varphi U_{\eta+i}\psi < t$ and $\varphi U_{\eta+j}\psi < t$, for $i < j$, and similarly for “ $>$ ”. Intuitively, sets that are not minimal hold redundant assertions, and can be ignored. Accordingly, we restrict the state space of the NBA to have only minimal sets.

Lemma 2. *For an $\text{LTL}^{\text{disc}}[\mathcal{D}]$ formula φ and $v \in [0, 1]$, the AWA $\mathcal{A}_{\varphi,v}$ constructed in Theorem 1 with state space Q can be translated to an NBA with $|Q|^{O(|\varphi|)}$ states.*

3.5 Decision Procedures for $\text{LTL}^{\text{disc}}[\mathcal{D}]$

Model checking and satisfiability. Consider a Kripke structure \mathcal{K} , an $\text{LTL}^{\text{disc}}[\mathcal{D}]$ formula φ , and a threshold v . By checking the emptiness of the intersection of \mathcal{K} with $\mathcal{A}_{\neg\varphi,1-v}$, we can solve the threshold model-checking problem. Indeed, $L(\mathcal{A}_{\neg\varphi,1-v}) \cap L(\mathcal{K}) \neq \emptyset$ iff there exists a lasso computation π that is induced by \mathcal{K} such that $\llbracket \pi, \varphi \rrbracket < v$, which happens iff it is not true that $\llbracket \mathcal{K}, \varphi \rrbracket \geq v$.

The complexity of the model-checking procedure depends on the discounting functions in \mathcal{D} . For the set of exponential-discounting functions E , we provide the following concrete complexities, showing that it stays in the same complexity classes of standard LTL model-checking.

Theorem 3. *For a Kripke structure \mathcal{K} , an $\text{LTL}^{\text{disc}}[E]$ formula φ , and a threshold $v \in [0, 1] \cap \mathbb{Q}$, the problem of deciding whether $\llbracket \mathcal{K}, \varphi \rrbracket > v$ is in NLOGSPACE in the number of states of \mathcal{K} , and in PSPACE in $|\langle \varphi \rangle|$ and in the description of v .*

Proof: By Theorem 2 and Lemma 2, the size of an NBA \mathcal{B} corresponding to φ and v is singly exponential in $|\langle \varphi \rangle|$ and in the description of v . Hence, we can check the emptiness of the intersection of \mathcal{K} and \mathcal{B} via standard “on the fly” procedures, getting the stated complexities. \square

Note that the complexity in Theorem 3 is only NLOGSPACE in the system, since our solution does not analyze the Kripke structure, but only takes its product with the specification’s automaton. This is in contrast to the approach of model checking temporal logic with (non-discounting) accumulative values, where, when decidable, involves a doubly-exponential dependency on the size of the system [4].

Finally, observe that the NBA obtained in Lemma 2 can be used to solve the threshold-satisfiability problem: given an $\text{LTL}^{\text{disc}}[\mathcal{D}]$ formula φ and a threshold $v \in [0, 1]$, we can decide whether there is a computation π such that $\llbracket \pi, \varphi \rrbracket \sim v$, for $\sim \in \{<, >\}$, and return such a computation when the answer is positive. This is done by simply deciding whether there exists a word that is accepted by the NBA.

Threshold synthesis Consider an $\text{LTL}^{\text{disc}}[\mathcal{D}]$ formula φ , and assume a partition of the atomic propositions in φ to input and output signals, we can use the NBA $\mathcal{A}_{\varphi,v}$ in order to address the *synthesis* problem, as stated in the following theorem (see Appendix A.6 for the proof).

Theorem 4. *Consider an $\text{LTL}^{\text{disc}}[\mathcal{D}]$ formula φ . If there exists a transducer \mathcal{T} all of whose computations π satisfy $\llbracket \pi, \varphi \rrbracket > v$, then we can generate a transducer \mathcal{T} all of whose computations τ satisfy $\llbracket \tau, \varphi \rrbracket \geq v$.*

4 Adding Propositional Quality Operators

As model checking is decidable for $\text{LTL}^{\text{disc}}[\mathcal{D}]$, one may wish to push the limit and extend the expressive power of the logic. In particular, of great interest is the combining of discounting with propositional quality operators [1].

4.1 Adding the Average Operator

A well-motivated extension is the introduction of the average operator \oplus , with the semantics $\llbracket \pi, \varphi \oplus \psi \rrbracket = \frac{\llbracket \pi, \varphi \rrbracket + \llbracket \pi, \psi \rrbracket}{2}$. The work in [1] proves that extending LTL by this operator, as well as with other propositional quantitative operators, enables clean specification of quality and results in a logic for which the model-checking problem can be solved in PSPACE.

We show that adding the \oplus operator to $\text{LTL}^{\text{disc}}[\mathcal{D}]$ gives a logic, denoted $\text{LTL}^{\text{disc}\oplus}[\mathcal{D}]$, for which the validity problem is undecidable. The validity problem asks, given an $\text{LTL}^{\text{disc}\oplus}[\mathcal{D}]$ formula φ over the atomic propositions AP and a threshold $v \in [0, 1]$, whether $\llbracket \pi, \varphi \rrbracket > v$ for every $\pi \in (2^{AP})^\omega$.

In the undecidability proof, we show a reduction from the 0-halting problem for two-counter machines. A *two-counter machine* \mathcal{M} is a sequence (l_1, \dots, l_n) of commands involving two counters x and y . We refer to $\{1, \dots, n\}$ as the *locations* of the machine. There are five possible forms of commands:

$$\text{INC}(c), \text{DEC}(c), \text{GOTO } l_i, \text{ IF } c=0 \text{ GOTO } l_i \text{ ELSE GOTO } l_j, \text{ HALT},$$

where $c \in \{x, y\}$ is a counter and $1 \leq i, j \leq n$ are locations. Since we can always check whether $c = 0$ before a $\text{DEC}(c)$ command, we assume that the machine never reaches $\text{DEC}(c)$ with $c = 0$. That is, the counters never have negative values. Given a counter machine \mathcal{M} , deciding whether \mathcal{M} halts is known to be undecidable [20]. Given \mathcal{M} , deciding whether \mathcal{M} halts with both counters having value 0, termed the *0-halting problem*, is also undecidable: given a counter machine \mathcal{M} , we can replace every HALT command with a code that clears the counters before halting. In fact, from this we see that the promise problem of deciding whether \mathcal{M} 0-halts given the promise that either it 0-halts, or it does not halt at all, is also undecidable.

Theorem 5. *The validity problem for $\text{LTL}^{\text{disc}\oplus}[\mathcal{D}]$ is undecidable (for every nonempty set of discounting functions \mathcal{D}).*

Proof: We start by showing a reduction from the 0-halting problem for two-counter machines to the following problem: given an $\text{LTL}^{\text{disc}\oplus}[\mathcal{D}]$ formula φ over the atomic propositions AP , whether there exists a path $\pi \in AP^\omega$ such that $\llbracket \pi, \varphi \rrbracket \geq \frac{1}{2}$. We dub this the $\frac{1}{2}$ -co-validity problem. We will later reduce this problem to the (complement of the) validity problem.

Let \mathcal{M} be a two-counter machine with commands (l_1, \dots, l_n) . A *halting run* of a two-counter machine with commands from the set $L = \{l_1, \dots, l_n\}$ is a sequence $\rho = \rho_1, \dots, \rho_m \in (L \times \mathbb{N} \times \mathbb{N})^*$ such that the following hold.

1. $\rho_1 = \langle l_1, 0, 0 \rangle$.
2. For all $1 < i \leq m$, let $\rho_{i-1} = \langle l_k, \alpha, \beta \rangle$ and $\rho_i = \langle l', \alpha', \beta' \rangle$. Then, the following hold.
 - If l_k is an $\text{INC}(x)$ command (resp. $\text{INC}(y)$), then $\alpha' = \alpha + 1, \beta' = \beta$ (resp. $\beta' = \beta + 1, \alpha' = \alpha$), and $l' = l_{k+1}$.
 - If l_k is a $\text{DEC}(x)$ command (resp. $\text{DEC}(y)$), then $\alpha' = \alpha - 1, \beta' = \beta$ (resp. $\beta' = \beta - 1, \alpha' = \alpha$), and $l' = l_{k+1}$.
 - If l_k is a $\text{GOTO } l_s$ command, then $\alpha' = \alpha, \beta' = \beta$, and $l' = l_s$.
 - If l_k is an $\text{IF } x=0 \text{ GOTO } l_s \text{ ELSE GOTO } l_t$ command, then $\alpha' = \alpha, \beta' = \beta$, and $l' = l_s$ if $\alpha = 0$, and $l' = l_t$ otherwise.
 - If l_k is an $\text{IF } y=0 \text{ GOTO } l_s \text{ ELSE GOTO } l_t$ command, then $\alpha' = \alpha, \beta' = \beta$, and $l' = l_s$ if $\beta = 0$, and $l' = l_t$ otherwise.
 - If l' is a HALT command, then $i = m$. That is, a run does not continue after HALT.
3. $\rho_m = \langle l_k, \alpha, \beta \rangle$ such that l_k is a HALT command.

Observe that the machine \mathcal{M} is deterministic. We say that \mathcal{M} 0-halts if there exists $l \in L$ that is a HALT command, such that the run of \mathcal{M} ends in $\langle l, 0, 0 \rangle$.

We say that a sequence of commands $\tau \in L^*$ fits a run ρ , if τ is the projection of ρ on its first component.

We construct from \mathcal{M} an $\text{LTL}^{\text{disc}\oplus}[\mathcal{D}]$ formula φ such that \mathcal{M} 0-halts iff there exists a computation π such that $\llbracket \pi, \varphi \rrbracket \geq \frac{1}{2}$. The idea behind the construction is as follows. φ reads computations over the atomic propositions $AP = \{1, \dots, n, \#, x, y\}$, where $1, \dots, n$ are the commands of \mathcal{M} . The computation that φ reads corresponds to a description of a run of \mathcal{M} , where every triplet $\langle l_i, \alpha, \beta \rangle$ is encoded as the string $ix^\alpha y^\beta \#$. We ensure that computations that satisfy φ with value greater than 0 are such that in every position only a single atomic proposition is true.

Example 2. Consider the following machine \mathcal{M} :

l_1 : INC(x)
 l_2 : IF $x=0$ GOTO l_6 ELSE GOTO l_3
 l_3 : INC(y)
 l_4 : DEC(x)
 l_5 : GOTO (l_2)
 l_6 : DEC(y)
 l_7 : HALT

The command sequence that represents the run of this machine is

$$\langle l_1, 0, 0 \rangle, \langle l_2, 1, 0 \rangle, \langle l_3, 1, 0 \rangle, \langle l_4, 1, 1 \rangle, \langle l_5, 0, 1 \rangle, \langle l_2, 0, 1 \rangle, \langle l_6, 0, 1 \rangle, \langle l_7, 0, 0 \rangle$$

and the encoding of it as a computation is

$$1\#2x\#3x\#4xy\#5y\#2y\#6y\#7\#$$

The formula φ “states” (recall that the setting is quantitative, not Boolean) the following properties of the computation π :

1. The first configuration in π is the initial configuration of \mathcal{M} ($\langle l_1, 0, 0 \rangle$, or $1\#$ in our encoding).
2. The last configuration in π is $\langle l, 0, 0 \rangle$ (or k in our encoding), where l can be any line whose command is HALT.
3. π represents a legal run of \mathcal{M} , up to the consistency of the counters between transitions.
4. The counters are updated correctly between configurations.

Properties 1-3 are can easily be specified by a LTL formulas, such that computations which satisfy properties 1-3 get satisfaction value 1. Property 4 utilizes the expressive power of $\text{LTL}^{\text{disc}\oplus}[\mathcal{D}]$, as we now demonstrate. The intuition behind property 4 is the following. We need to compare the value of a counter before and after a command, such that the formula takes a low value if a violation is encountered, and a high value otherwise. Specifically, the formula we construct takes value $\frac{1}{2}$ if no violation occurred, and a lower value if a violation did occur.

We start with a simpler case, to demonstrate the point. Let $\eta \in \mathcal{D}$ be a discounting function. Consider the formula $\text{Count}A := a \cup_{\eta} \neg a$ and the computation $a^i b^j \#^\omega$. It holds that $\llbracket a^i b^j, \text{Count}A \rrbracket = \eta(i)$. Similarly, it holds that $\llbracket a^i b^j \#^\omega, a \cup (b \cup_{\eta} \neg b) \rrbracket = \eta(j)$. Denote the latter by $\text{Count}B$. Let

$$\text{Compare}AB := (\text{Count}A \oplus \neg \text{Count}B) \wedge (\neg \text{Count}A \oplus \text{Count}B).$$

We now have that

$$\llbracket a^i b^j \#^\omega, \text{Compare}AB \rrbracket = \min \left\{ \frac{\eta(i) + 1 - \eta(j)}{2}, \frac{\eta(j) + 1 - \eta(i)}{2} \right\} = \frac{1}{2} - \frac{|\eta(i) - \eta(j)|}{2}$$

and observe that the latter is $\frac{1}{2}$ iff $i = j$, and is less than $\frac{1}{2}$ otherwise. This is because η is strictly decreasing, and in particular an injection.

Thus, we can compare counters. To apply this technique to the encoding of a computation, we only need some technical formulas to “parse” the input and find consecutive occurrences of a counter.

We now dive into the technical definition of φ . The atomic propositions are $AP = \{1, \dots, n, \#, x, y\}$ (where l_1, \dots, l_n are the commands of \mathcal{M}). We let $\varphi := \text{CheckCmds} \wedge \text{CheckInit} \wedge \text{CheckFinal} \wedge \text{CheckCounters} \wedge \text{ForceSingletons}$

ForceSingletons: This formula ensures that for a computation to get a value of more than 0, every letter in the computation must be a singleton. Formally,

$$\text{ForceSingletons} := \mathbf{G} \left(\bigvee_{p \in AP} (p \wedge \bigwedge_{q \in AP \setminus \{p\}} \neg q) \right).$$

CheckInit and CheckFinal: These formulas check that the initial and final configurations are correct, and that after the final configuration, there are only $\#$ s.

$$\text{CheckInit} := 1 \wedge \mathbf{X}\#.$$

Let $I = \{i : l_i = \text{HALT}\}$, we define

$$CheckFinal := G((\bigvee_{i \in I} i) \rightarrow XG\#).$$

Note that $CheckFinal$ also ensures that there counters are 0.

CheckCmds: This formula verifies that the local transitions follow the instructions in the counter machine, ignoring the consistency of the counter values, but enforcing that a jump behaves according to the counters. We start by defining, for every $i \in \{1, \dots, n\}$, the formula:

$$waitfor(i) := (x \vee y)U(\# \wedge Xi)$$

Intuitively, a computation satisfies this formula (i.e., gets value 1) iff it reads counter descriptions until the next delimiter, and the next command is l_i .

Now, for every $i \in \{1, \dots, n\}$ we define ψ_i as follows.

- If $l_i = \text{GOTO } l_j$, then $\psi_i := Xwaitfor(j)$.
- If $l_i \in \{\text{INC}(c), \text{DEC}(c) : c \in \{x, y\}\}$, then $\psi_i := Xwaitfor(i + 1)$.⁴
- If $l_i = \text{IF } x=0 \text{ GOTO } l_j \text{ ELSE GOTO } l_k$ then $\psi_i := X((x \rightarrow waitfor(k)) \wedge ((\neg x) \rightarrow waitfor(j)))$.
- If $l_i = \text{IF } y=0 \text{ GOTO } l_j \text{ ELSE GOTO } l_k$ then $\psi_i := X(((xUy) \wedge waitfor(k)) \vee ((xU\#) \wedge waitfor(j)))$.
- If $l_i = \text{HALT}$ we do not really need additional constraints, due to $CheckFinal$. Thus we have $\psi_i = \text{True}$.

Finally, we define $CheckCmds := G \bigwedge_{i \in \{1, \dots, n\}} (i \rightarrow \psi_i)$.

CheckCounters: This is the heart of the construction. The formula checks whether consecutive occurrences of the counters match the transition between the commands. We start by defining $countX := xU_\eta \neg x$ and $countY := xU(yU_\eta \neg y)$. Similarly, we have $countX^{-1} = xU_\eta X \neg x$ and $countY^{-1} = xU(yU_\eta X \neg y)$.

We need to define a formula to handle some edge cases.

Let $I_{\text{HALT}} = \{i : l_i = \text{HALT}\}$, and similarly define $I_{\text{DEC}(x)}$ and $I_{\text{DEC}(y)}$. We define

$$\begin{aligned} Last := & \bigvee_{i \in I_{\text{DEC}(x)}} i \wedge X(x \wedge X(\# \wedge X \bigvee_{k \in I_{\text{HALT}}} k)) \vee \\ & \bigvee_{i \in I_{\text{DEC}(y)}} i \wedge X(y \wedge X(\# \wedge X \bigvee_{k \in I_{\text{HALT}}} k)) \vee \\ & \bigvee_{i \in \{1, \dots, n\}} i \wedge X(\# \wedge X \bigvee_{k \in I_{\text{HALT}}} k) \end{aligned}$$

Intuitively, $Last$ holds exactly in the last transition, that is - before the final 0-halting configuration.

Testing the counters involves six types of comparisons: checking equality, increase by 1, and decrease by 1 for each of the two counters. We define the formulas below for these tests. To explain the formulas, consider for example the formula $Comp(x, =)$. This formula compares the number of x 's in the current configuration, with the number of x 's in the next configuration. The comparison is based on the comparison we explained above, and is augmented by some parsing, as we need to reach the next configuration before comparing.

- $Comp(x, =) := (countX \oplus (xU(yU(\# \wedge XX \neg countX)))) \wedge ((\neg countX) \oplus (xU(yU(\# \wedge XX countX))))$.
- $Comp(y, =) := (countY \oplus (xU(yU(\# \wedge XX \neg countY)))) \wedge ((\neg countY) \oplus (xU(yU(\# \wedge XX countY))))$.
- $Comp(x, +1) := (countX \oplus (xU(yU(\# \wedge XX \neg countX^{-1})))) \wedge ((\neg countX) \oplus (xU(yU(\# \wedge XX countX^{-1}))))$.
- $Comp(y, +1) := (countY \oplus (xU(yU(\# \wedge XX \neg countY^{-1})))) \wedge ((\neg countY) \oplus (xU(yU(\# \wedge XX countY^{-1}))))$.
- $Comp(x, -1) := (countX^{-1} \oplus (xU(yU(\# \wedge XX \neg countX)))) \wedge ((\neg countX^{-1}) \oplus (xU(yU(\# \wedge XX countX))))$.

⁴ if $i = n$ then this line can be omitted from the initial machine, so w.l.o.g this does not happen.

- $Comp(y, -1) := (countY^{-1} \oplus (xU(yU(\# \wedge XX\neg countY))))$
 $\wedge ((\neg countY^{-1}) \oplus (xU(yU(\# \wedge XXcountY))))$

Now, for every $i \in \{1, \dots, n\}$ we define ξ_i as follows.

- If $l_i \in \{\text{GOTO } l_j, \text{IF } c=0 \text{ GOTO } l_j \text{ ELSE GOTO } l_k : c \in \{x, y\}\}$, we need to make sure the value of the counters do not change. We define
 $\xi_i := (X(comp(x, =) \wedge comp(y, =)) \vee Last.$
- If $l_i = \text{INC}(x)$, we need to make sure that x increases and y does not change. We define
 $\xi_i := (X(comp(x, +1) \wedge comp(y, =)) \vee Last.$
- If $l_i = \text{INC}(y)$, we define
 $\xi_i := (X(comp(x, =) \wedge comp(y, +1)) \vee Last.$
- If $l_i = \text{DEC}(x)$, we define
 $\xi_i := (X(comp(x, -1) \wedge comp(y, =)) \vee Last.$
- If $l_i = \text{DEC}(y)$, we define
 $\xi_i := (X(comp(x, =) \wedge comp(y, -1)) \vee Last.$
- If $l_i = \text{HALT}$ we do not need additional constraints, due to *CheckFinal*. Thus we have $\psi_i = \text{True}$.

Finally, we define $CheckCounters := G \bigwedge_{i \in \{1, \dots, n\}} (i \rightarrow \xi_i)$

The correctness of the construction is obvious, once one verifies that the defined formulas indeed test what they claim to. Thus, we conclude that \mathcal{M} 0-halts iff there exists a computation π such that $\llbracket \pi, \varphi \rrbracket \geq \frac{1}{2}$.

Finally, we reduce the $\frac{1}{2}$ -co-validity problem to the complement of the validity problem: given a formula φ , the reduction outputs $\langle \neg\varphi, \frac{1}{2} \rangle$. Now, there exists a computation π such that $\llbracket \pi, \varphi \rrbracket \geq \frac{1}{2}$ iff there exists a computation π such that $\llbracket \pi, \neg\varphi \rrbracket \leq \frac{1}{2}$, iff it is not true that $\llbracket \pi, \neg\varphi \rrbracket > \frac{1}{2}$ for every computation π . Thus, φ is $\frac{1}{2}$ -co-valid iff $\neg\varphi$ is not valid for threshold $\frac{1}{2}$. We conclude that the validity problem is undecidable. \square

Studying the proof of Theorem 5, we can actually formulate the reduction more carefully as follows.

Lemma 3. *Given a two-counter machine \mathcal{M} that is promised to either 0-halt, or not to halt at all, there exists an $LTL^{disc\oplus}[\mathcal{D}]$ formula φ such that for every computation π that represents a computation of \mathcal{M} , the following hold.*

1. *If π is a legal halting computation of \mathcal{M} , then $\llbracket \pi, \varphi \rrbracket = \frac{1}{2}$.*
2. *If π cheats in a transition between commands, then $\llbracket \pi, \varphi \rrbracket = 1$.*
3. *If π cheats in the counter values, then $\llbracket \pi, \varphi \rrbracket = \frac{1}{2} + \epsilon$ such that $\epsilon \geq \frac{1}{2}(\eta(i) - \eta(i+1))$ for the minimal difference $\eta(i) - \eta(i+1)$ where i is a counter value in π .*

We now turn to show that the *strict* model-checking problem and the *strict* model-checking problem are undecidable for $LTL^{disc\oplus}[\mathcal{D}]$ as well. The strict model-checking problem is to decide, given a Kripke structure \mathcal{K} , a formula φ , and a threshold v , whether $\llbracket \mathcal{K}, \varphi \rrbracket > v$.

Theorem 6. *The strict model-checking problem for $LTL^{disc\oplus}[\mathcal{D}]$ is undecidable (for every nonempty set of Discounting functions \mathcal{D}).*

Proof: Assume by way of contradiction that the strict model-checking problem is decidable. We show how to decide the 0-halting promise problem for two-counter machines, thus reaching a contradiction.

Given a two-counter machine \mathcal{M} that is promised to either 0-halt, or not halt at all, construct the formula φ as per Lemma 3, and consider the Kripke structure \mathcal{K} that generates every computation. Observe that by Lemma 3 it holds that $\llbracket \mathcal{K}, \varphi \rrbracket \geq \frac{1}{2}$.

Decide whether $\llbracket \mathcal{K}, \varphi \rrbracket > \frac{1}{2}$. If $\llbracket \mathcal{K}, \varphi \rrbracket > \frac{1}{2}$, then for every computation π it holds that $\llbracket \pi, \varphi \rrbracket > \frac{1}{2}$, and by Lemma 3 we conclude that \mathcal{M} does not halt.

If $\llbracket \mathcal{K}, \varphi \rrbracket \leq \frac{1}{2}$, then $\llbracket \mathcal{K}, \varphi \rrbracket = \frac{1}{2}$. We observe that there are now two possible cases:

1. \mathcal{M} halts.

2. \mathcal{M} does not halt, and for every n , there are computations that reach HALT while cheating in counter values larger than n , and not cheating in the commands.

We show how to distinguish between cases 1 and 2.

Consider the $\text{LTL}^{\text{disc} \oplus}[\mathcal{D}]$ formula

$$\psi = \text{True} \oplus (\text{G}(x \text{U}_\eta \neg x) \wedge \text{G}(y \text{U}_\eta \neg y)).$$

It is not hard to verify that for every computation π that represents a computation of \mathcal{M} , it holds that $\llbracket \pi, \psi \rrbracket = \frac{1}{2} + \frac{1}{2}\epsilon$, where $\epsilon = \inf(\eta(i) : i \text{ is the value of a counter in } \pi)$. Let $\xi = \varphi \vee \psi$.

If \mathcal{M} halts (case 1), then for every computation π we have one of the following.

- a. π describes a legal halting run of \mathcal{M} , in which case $\llbracket \pi, \varphi \rrbracket = \frac{1}{2}$ and $\llbracket \pi, \psi \rrbracket > \frac{1}{2} + \epsilon$ for some $\epsilon > 0$ (independent of π), since the counters are bounded. Thus, $\llbracket \pi, \xi \rrbracket > \frac{1}{2} + \epsilon$.
- b. π cheats in the commands, in which case $\llbracket \pi, \varphi \rrbracket = 1$, so $\llbracket \pi, \xi \rrbracket = 1$.
- c. π cheats in the counters, in which case, since the counters are bounded, the first cheat must occur with small counters, and thus $\llbracket \pi, \varphi \rrbracket > \frac{1}{2} + \epsilon$ for some $\epsilon > 0$ independent of π . So $\llbracket \pi, \xi \rrbracket > \frac{1}{2} + \epsilon$.

In all three cases, we get that $\llbracket \pi, \xi \rrbracket > \frac{1}{2} + \epsilon$, so $\llbracket \mathcal{K}, \xi \rrbracket > \frac{1}{2}$.

If \mathcal{M} does not halt (case 2), then for every n and for every computation π that cheats with counters larger than n , it holds that $\llbracket \pi, \varphi \rrbracket < \frac{1}{2} + \epsilon$ where $\epsilon = \epsilon(n) \rightarrow 0$ as $n \rightarrow \infty$, and since the counters in π are large, it also holds that $\llbracket \pi, \psi \rrbracket < \frac{1}{2} + \epsilon'$, where $\epsilon' = \epsilon'(n) \rightarrow 0$ as $n \rightarrow \infty$. We conclude that there exists a sequence of computations whose satisfaction values in ξ tend to $\frac{1}{2}$, and thus $\llbracket \mathcal{K}, \xi \rrbracket = \frac{1}{2}$.

Thus, in order to distinguish between cases 1 and 2, it is enough to decide whether $\llbracket \mathcal{K}, \xi \rrbracket > \frac{1}{2}$.

To conclude, the algorithm for deciding whether \mathcal{M} 0-halts is as follows. Start by constructing φ . If $\llbracket \mathcal{K}, \varphi \rrbracket > \frac{1}{2}$, then \mathcal{M} does not halt. Otherwise, construct ξ . If $\llbracket \mathcal{K}, \xi \rrbracket > \frac{1}{2}$, then \mathcal{M} halts, and otherwise \mathcal{M} does not halt. \square

Theorem 7. *The model-checking problem for $\text{LTL}^{\text{disc} \oplus}[\mathcal{D}]$ is undecidable (for every nonempty set of Discounting functions \mathcal{D}).*

Proof: Recall that for every Kripke structure \mathcal{K} and formula φ it holds that $\llbracket \mathcal{K}, \varphi \rrbracket \geq v$ iff there does not exist a computation π of \mathcal{K} such that $\llbracket \pi, \neg\varphi \rrbracket > 1 - v$.

We show that the latter problem is undecidable, even if we fix \mathcal{K} to be the system that generates every computation.

We show a reduction from the 0-halting promise problem to the latter problem. Given a two-counter machine \mathcal{M} that is promised to either 0-halt, or not halt at all, construct the formula φ as per Lemma 3 and the formula ψ such that for every computation π we have that $\llbracket \pi, \psi \rrbracket = \frac{1}{2} + \frac{1}{2}\epsilon$, where $\epsilon = \inf(\eta(i) - \eta(i+1) : i \text{ is the value of a counter in } \pi)$. The formula ψ can be defined as

$$\psi = \text{True} \oplus (\text{G}((x \vee \text{X}x) \text{U}_\eta (\neg x \wedge \neg \text{X}x)) \wedge (y \vee \text{X}y) \text{U}_\eta (\neg y \wedge \neg \text{X}y))$$

Let $\theta = (\neg\varphi) \oplus \psi$. We claim that \mathcal{M} halts iff there exists a computation π such that $\llbracket \pi, \theta \rrbracket > \frac{1}{2}$.

If \mathcal{M} halts, then for the computation π that describes the halting run of \mathcal{M} it holds that $\llbracket \pi, \varphi \rrbracket = \frac{1}{2}$, and thus $\llbracket \pi, \neg\varphi \rrbracket = \frac{1}{2}$. Since the counters in π are bounded (as the run is halting), then $\llbracket \pi, \psi \rrbracket > \frac{1}{2}$, and thus $\llbracket \pi, \theta \rrbracket > \frac{1}{2}$.

If \mathcal{M} does not halt, consider a computation π .

- If π cheats in the commands, then $\llbracket \pi, \neg\varphi \rrbracket = 0$, so $\llbracket \pi, \theta \rrbracket = 0 + \frac{1}{2}\llbracket \pi, \psi \rrbracket \leq \frac{1}{2}$.
- If π cheats in the counters, then $\llbracket \pi, \neg\varphi \rrbracket = \frac{1}{2} - \frac{1}{2}\epsilon$ and $\llbracket \pi, \psi \rrbracket = \frac{1}{2} + \frac{1}{2}\epsilon'$, where $\epsilon \geq \frac{1}{2}(\eta(i) - \eta(i+1)) = \epsilon'$ for the smallest difference $\eta(i) - \eta(i+1)$ in π . Thus, $\llbracket \pi, \theta \rrbracket \leq \frac{1}{2}$.

\square

4.2 Adding Unary Multiplication Operators

As we have seen in Section 4.1, adding the operator \oplus to $\text{LTL}^{\text{disc}}[\mathcal{D}]$ makes model checking undecidable. One may still want to find propositional quality operators that we can add to the logic preserving its decidability. In this section we describe one such operator. We extend $\text{LTL}^{\text{disc}}[\mathcal{D}]$ with the operator ∇_λ , for $\lambda \in (0, 1)$, with the semantics

$\llbracket \pi, \nabla_\lambda \varphi \rrbracket = \lambda \cdot \llbracket \pi, \varphi \rrbracket$. This operator allows the specifier to manually change the satisfaction value of certain subformulas. This can be used to express importance, reliability, etc. of subformulas. For example, in $G(\text{request} \rightarrow (\text{response} \vee \nabla_{\frac{2}{3}} X \text{response}))$, we limit the satisfaction value of computations in which a response is given with a delay to $\frac{2}{3}$.

Note that the operator ∇_λ is similar to a one-time application of $U_{\exp_\lambda^{+1}}$, thus $\nabla_\lambda \varphi$ is equivalent to $\text{False} U_{\exp_\lambda^{+1}} \psi$. In practice, it is better to handle ∇_λ formulas directly, by adding the following transitions to the construction in the proof of Theorem 1.

$$\delta(\nabla_\lambda \varphi > t, \sigma) = \begin{cases} \delta(\varphi > \frac{t}{\lambda}, \sigma) & \text{if } \frac{t}{\lambda} < 1, \\ \text{False} & \text{if } \frac{t}{\lambda} \geq 1, \end{cases} \quad \delta(\nabla_\lambda \varphi < t, \sigma) = \begin{cases} \delta(\varphi < \frac{t}{\lambda}, \sigma) & \text{if } \frac{t}{\lambda} \leq 1, \\ \text{True} & \text{if } \frac{t}{\lambda} > 1. \end{cases}$$

5 Extensions

5.1 LTL^{disc}[D] with Past Operators

One of the well-known augmentations of LTL is the addition of *past operators* [18]. These operators enable the specification of exponentially more succinct formulas, while preserving the PSPACE complexity of model checking. In this section, we add *discounting-past* operators to LTL^{disc}[D], and show how to perform model-checking on the obtained logic.

We add the operators $Y\varphi$, $\varphi S\psi$, and $\varphi S_\eta \varphi$ (for $\eta \in \mathcal{D}$) to LTL^{disc}[D], and denote the extended logic PLTL^{disc}[D], with the following semantics. For PLTL^{disc}[D] formulas φ, ψ , a function $\eta \in \mathcal{D}$, a computation π , and an index $i \in \mathbb{N}$, we have

- $\llbracket \pi^i, Y\varphi \rrbracket = \llbracket \pi^{i-1}, \varphi \rrbracket$ if $i > 0$, and 0 otherwise.
- $\llbracket \pi^i, \varphi S\psi \rrbracket = \max_{0 \leq j \leq i} \{ \min \{ \llbracket \pi^j, \psi \rrbracket, \min_{j < k \leq i} \{ \llbracket \pi^k, \varphi \rrbracket \} \} \}$.
- $\llbracket \pi^i, \varphi S_\eta \psi \rrbracket = \max_{0 \leq j \leq i} \{ \min \{ \eta(i-j) \llbracket \pi^j, \psi \rrbracket, \min_{j < k \leq i} \{ \eta(i-k) \llbracket \pi^k, \varphi \rrbracket \} \} \}$.

Observe that since the past is finite, then the semantics for past operators can use min and max instead of inf and sup.

As in LTL^{disc}[D], our solution for the PLTL^{disc}[D] model-checking problem is by translating PLTL^{disc}[D] formulas to automata. The construction extends the construction for the Boolean case, which uses 2-way weak alternating automata (2AWA). The details of the construction appear in Appendix A.7. The use of the obtained automata in decision procedures is similar to that in Section 3. In particular, it follows that the model-checking problem for PLTL^{disc}[D] with exponential discounting is in PSPACE.

5.2 Weighted Systems

A central property of the logic LTL^{disc}[D] is that the verified system need not be weighted in order to get a quantitative satisfaction – it stems from taking into account the delays in satisfying the requirements. Nevertheless, LTL^{disc}[D] also naturally fits weighted systems, where the atomic propositions have a value between 0 and 1.

A *weighted Kripke structure* is a tuple $\mathcal{K} = \langle AP, S, I, \rho, L \rangle$, where AP, S, I , and ρ are as in Boolean Kripke structures, and $L : S \rightarrow [0, 1]^{AP}$ maps each state to a weighted assignment to the atomic propositions. Thus, the value $L(s)(p)$ of an atomic proposition $p \in AP$ in a state $s \in S$ is a value in $[0, 1]$. The semantics of LTL^{disc}[D] with respect to a weighted computation coincides with the one for non-weighted systems, except that for an atomic proposition p , we have that $\llbracket \pi, p \rrbracket = L(\pi_0)(p)$.

It is possible to extend the construction of $\mathcal{A}_{\varphi, v}$ described in Section 3.2 to an alphabet W^{AP} , where W is a set of possible values for the atomic propositions. Indeed, we only have to adjust the transition for states that correspond to atomic propositions, as follows: for $p \in AP$, $v \in [0, 1]$, and $\sigma \in W^{AP}$, we have that

$$-\delta(p > v, \sigma) = \begin{cases} \text{True} & \text{if } \sigma(p) > v, \\ \text{False} & \text{otherwise.} \end{cases} \quad -\delta(p < v, \sigma) = \begin{cases} \text{True} & \text{if } \sigma(p) < v, \\ \text{False} & \text{otherwise.} \end{cases}$$

5.3 Changing the Tendency of Discounting

One may observe that in our discounting scheme, the value of future formulas is discounted toward 0. This, in a way, reflects an intuition that we are pessimistic about the future, or at least we are impatient. While in some cases this fits

the needs of the specifier, it may well be the case that we are ambivalent to the future. To capture this notion, one may want the discounting to tend to $\frac{1}{2}$. Other values are also possible. For example, it may be that we are optimistic about the future, say when a system improves its performance while running and we know that components are likely to function better in the future. We may then want the discounting to tend, say, to $\frac{3}{4}$.

To capture this notion, we define the operator $O_{\eta,z}$, parameterized by $\eta \in \mathcal{D}$ and $z \in [0, 1]$, with the following semantics. $\llbracket \pi, \varphi O_{\eta,z} \psi \rrbracket = \sup_{i \geq 0} \{ \min \{ \eta(i) \llbracket \pi^i, \psi \rrbracket + (1 - \eta(i))z, \min_{0 \leq j < i} \eta(j) \llbracket \pi^j, \varphi \rrbracket + (1 - \eta(j))z \} \}$. The discounting function η determines the rate of convergence, and z determines the limit of the discounting. The longer it takes to fulfill the “eventuality”, the closer the satisfaction value gets to z . We observe that $\varphi U_{\eta} \psi \equiv \varphi O_{\eta,0} \psi$.

Example 3. Consider a process scheduler. The scheduler decides which process to run at any given time. The scheduler may also run a defragment tool, but only if it is not in expense of other processes. This can be captured by the formula $\varphi = \text{True} O_{\eta, \frac{1}{2}} \text{defrag}$. Thus, the defragment tool is a “bonus”: if it runs, then the satisfaction value is above $\frac{1}{2}$, but if it does not run, the satisfaction value is $\frac{1}{2}$. Treating 1 as “good” and 0 as “bad” means that $\frac{1}{2}$ is ambivalent.

We claim that Theorem 1 holds under the extension of LTL^{disc} $[\mathcal{D}]$ with the operator O . Indeed, the construction of the AWA is augmented as follows. For $\varphi = \psi_1 O_{\eta,z} \psi_2$, denote $\frac{t - (1 - \eta(0))z}{\eta(0)}$ by τ . One may observe that the conditions on $\frac{t}{\eta(0)}$ correspond to conditions on τ when dealing with O . Accordingly, the transitions from the state $(\psi_1 O_{\eta,z} \psi_2 > t)$ are defined as follows.

First, if $t = z$, then $\tau = t$ and we identify the state $(\psi_1 O_{\eta,z} \psi_2 > t)$ with the state $(\psi_1 U \psi_2 > t)$. Otherwise, $z \neq t$ and we define:

$$\begin{aligned} - \delta((\psi_1 O_{\eta,z} \psi_2 > t), \sigma) &= \begin{cases} \delta((\psi_2 > \tau), \sigma) \vee \\ [\delta((\psi_1 > \tau), \sigma) \wedge (\psi_1 O_{\eta+1,z} \psi_2 > t)] & \text{if } 0 \leq \tau < 1, \\ \text{False} & \text{if } \tau \geq 1, \\ \text{True} & \text{if } \tau < 0. \end{cases} \\ - \delta((\psi_1 O_{\eta,z} \psi_2 < t), \sigma) &= \begin{cases} \delta((\psi_2 < \tau), \sigma) \wedge \\ [\delta((\psi_1 < \tau), \sigma) \vee (\psi_1 O_{\eta+1,z} \psi_2 < t)] & \text{if } 0 < \tau \leq 1, \\ \text{True} & \text{if } \tau > 1, \\ \text{False} & \text{if } \tau \leq 0. \end{cases} \end{aligned}$$

The correctness of the construction is proved in Appendix A.8.

6 Discussion

An ability to specify and to reason about quality would take formal methods a significant step forward. Quality has many aspects, some of which are propositional, such as prioritizing one satisfaction scheme on top of another, and some are temporal, for example having higher quality for implementations with shorter delays. In this work we provided a solution for specifying and reasoning about temporal quality, augmenting the commonly used linear temporal logic (LTL). A satisfaction scheme, such as ours, that is based on elapsed times introduces a big challenge, as it implies infinitely many satisfaction values. Nonetheless, we showed the decidability of the model-checking problem, and for the natural exponential-decaying satisfactions, the complexity remains as the one for standard LTL, suggesting the interesting potential of the new scheme. As for combining propositional and temporal quality operators, we showed that the problem is, in general, undecidable, while certain combinations, such as adding priorities, preserve the decidability and the complexity.

Acknowledgement. We thank Eleni Mandrali for pointing to an error in an earlier version of the paper.

References

1. S. Almagor, U. Boker, and O. Kupferman. Formalizing and reasoning about quality. In *40th ICALP*, 2013.
2. S. Almagor, Y. Hirshfeld, and O. Kupferman. Promptness in ω -regular automata. In *8th ATVA, LNCS 6252*, pages 22–36, 2010.
3. M. Bojańczyk and T. Colcombet. Bounds in ω -regularity. In *21st LICS*, pages 285–296, 2006.

4. U. Boker, K. Chatterjee, T.A. Henzinger, and O. Kupferman. Temporal Specifications with Accumulative Values. In *26th LICS*, pages 43–52, 2011.
5. E. Clarke, O. Grumberg, and D. Peled. *Model Checking*. MIT Press, 1999.
6. M. Dam. CTL* and ECTL* as fragments of the modal μ -calculus. *TCS*, 126:77–96, 1994.
7. L. de Alfaro, M. Faella, T. Henzinger, R. Majumdar, and M. Stoelinga. Model checking discounted temporal properties. *TCS*, 345(1):139–170, 2005.
8. L. de Alfaro, T. Henzinger, and R. Majumdar. Discounting the future in systems theory. In *30th ICALP, LNCS 2719*, pages 1022–1037, 2003.
9. M. Droste and G. Rahonis. Weighted automata and weighted logics with discounting. *TCS*, 410(37):3481–3494, 2009.
10. M. Droste and H. Vogler. Weighted automata and multi-valued logics over arbitrary bounded lattices. *TCS*, 418:14–36, 2012.
11. M. Faella, A. Legay, and M. Stoelinga. Model checking quantitative linear time logic. *Electr. Notes Theor. Comput. Sci.*, 220(3):61–77, 2008.
12. P. Gastin and D. Oddoux. Fast LTL to Büchi automata translation. In *13th CAV, LNCS 2102*, pages 53–65. Springer, 2001.
13. S.H. Kan. Metrics and Models in Software Quality Engineering. *Addison-Wesley Longman Publishing Co.*, 2002.
14. O. Kupferman, N. Piterman, and M.Y. Vardi. From Liveness to Promptness. In *19th CAV, LNCS 4590*, pages 406–419. Springer, 2007.
15. D. Krob. The equality problem for rational series with multiplicities in the tropical semiring is undecidable. *International Journal of Algebra and Computation*, 4(3):405–425, 1994.
16. M. Kwiatkowska. Quantitative verification: models techniques and tools. In *ESEC/SIGSOFT FSE*, pages 449–458, 2007.
17. F. Laroussinie and P. Schnoebelen. A hierarchy of temporal logics with past. In *11th STACS*, pages 47–58, 1994.
18. O. Lichtenstein, A. Pnueli, and L. Zuck. The glory of the past. In *Logics of Programs, LNCS 193*, pages 196–218. Springer, 1985.
19. E. Mandrali. Weighted LTL with discounting. In *CIAA, LNCS 7381*, pages 353–360, 2012.
20. M. Minsky. *Computation: Finite and Infinite Machines*. Prentice Hall, 1 edition, 1967.
21. S. Miyano and T. Hayashi. Alternating finite automata on ω -words. *TCS*, 32:321–330, 1984.
22. M. Mohri. Finite-state transducers in language and speech processing. *Computational Linguistics*, 23(2):269–311, 1997.
23. S. Moon, K. Lee, and D. Lee. Fuzzy branching temporal logic. *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, 34(2):1045–1055, 2004.
24. A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *Proc. 16th POPL*, pages 179–190, 1989.
25. L. Shapley. Stochastic games. In *Proc. of the National Academy of Science*, vol. 39, 1953.
26. D. Spinellis. Code Quality: The Open Source Perspective. *Addison-Wesley Professional*, 2006.
27. M.Y. Vardi. An automata-theoretic approach to linear temporal logic. In *Logics for Concurrency: Structure versus Automata, LNCS 1043*, pages 238–266, 1996.
28. M. Vardi and P. Wolper. An automata-theoretic approach to automatic program verification. In *1st LICS*, pages 332–344, 1986.

A Proofs

A.1 Proof of Lemma 1

We construct φ^+ and $\varphi^{<1}$ by induction on the structure of φ as follows. In all cases but the U case we do not use the assumption that $\pi = u \cdot v^\omega$ and prove an “iff” criterion.

- If φ is of the form **True**, **False**, or p , for an atomic proposition p , then $\varphi^+ = \varphi$ and $\varphi^{<1} = \neg\varphi$. Correctness is trivial.
- If φ is of the form $\psi_1 \vee \psi_2$, then $\varphi^+ = \psi_1^+ \vee \psi_2^+$ and $\varphi^{<1} = \psi_1^{<1} \wedge \psi_2^{<1}$. Indeed, for every computation π we have that $\llbracket \pi, \varphi \rrbracket > 0$ iff either $\llbracket \pi, \psi_1 \rrbracket > 0$ or $\llbracket \pi, \psi_2 \rrbracket > 0$, and $\llbracket \pi, \varphi \rrbracket < 1$ iff both $\llbracket \pi, \psi_1 \rrbracket < 1$ and $\llbracket \pi, \psi_2 \rrbracket < 1$.
- If φ is of the form $X\psi_1$, then $\varphi^+ = X(\psi_1^+)$ and $\varphi^{<1} = X(\psi_1^{<1})$. Correctness is trivial.
- If φ is of the form $\psi_1 U \psi_2$, then $\varphi^+ = \psi_1^+ U \psi_2^+$ and $\varphi^{<1} = \neg((\neg(\psi_1^{<1})) U (\neg(\psi_2^{<1})))$.

We start with φ^+ . For every computation π we have that $\llbracket \pi, \varphi \rrbracket > 0$ iff there exists $i \geq 0$ such that $\llbracket \pi^i, \psi_2 \rrbracket > 0$ and for every $0 \leq j < i$ it holds that $\llbracket \pi^j, \psi_1 \rrbracket > 0$. This happens iff π satisfies $\psi_1^+ U \psi_2^+$.

Before we turn to the case of $\varphi^{<1}$, let us note that readers familiar with the release (R) operator of LTL may find it clearer to observe that $\varphi^{<1} = \psi_1^{<1} R \psi_2^{<1}$, which perhaps gives a clearer intuition for the correctness of the construction.

Now, if $\llbracket \pi, \varphi \rrbracket < 1$, then for every $i \geq 0$ it holds that either $\llbracket \pi^i, \psi_2 \rrbracket < 1$ or $\llbracket \pi^j, \psi_1 \rrbracket < 1$ for some $0 \leq j < i$. Thus, for every $i \geq 0$, either $\pi^i \models \psi_2^{<1}$, or $\pi^j \models \psi_1^{<1}$ for some $0 \leq j < i$. So for every $i \geq 0$, either $\pi^i \not\models \neg(\psi_2^{<1})$, or $\pi^j \not\models \neg(\psi_1^{<1})$ for some $0 \leq j < i$. It follows that $\pi \not\models (\neg(\psi_1^{<1})) U (\neg(\psi_2^{<1}))$. Equivalently, $\pi \models \neg((\neg(\psi_1^{<1})) U (\neg(\psi_2^{<1})))$.

Conversely, if $\pi = u \cdot v^\omega$ and $\pi \models \neg((\neg(\psi_1^{<1})) U (\neg(\psi_2^{<1})))$, then $\pi \not\models (\neg(\psi_1^{<1})) U (\neg(\psi_2^{<1}))$, so for every $i \geq 0$, either $\pi^i \models \psi_2^{<1}$ or $\pi^j \models \psi_1^{<1}$ for some $0 \leq j < i$. By the induction hypothesis, for every $i \geq 0$, either $\llbracket \pi^i, \psi_2 \rrbracket < 1$ or $\llbracket \pi^j, \psi_1 \rrbracket < 1$ for some $0 \leq j < i$. We now use the assumption that $\pi = u \cdot v^\omega$ to observe that the sup in the expression for $\llbracket \pi, \varphi \rrbracket$ is attained as a max, as there are only finitely many distinct suffixes for π (namely $\pi^0, \dots, \pi^{|u|+|v|-1}$). Thus, since all the elements in the max are strictly smaller than 1, we conclude that $\llbracket \pi, \varphi \rrbracket < 1$.

- If $\varphi = \neg\psi$, then $\varphi^+ = \psi^{<1}$ and $\varphi^{<1} = \psi^+$. Again, correctness is trivial.
- If $\varphi = \psi_1 U_\eta \psi_2$ for $\eta \in \mathcal{D}$, then $\varphi^+ = \psi_1^+ U \psi_2^+$. Indeed, since $\eta(i) > 0$ for all $i \geq 0$, then $\varphi^+ = \psi_1 U \psi_2^+$. Now, $\varphi^{<1}$ is defined as follows. First, if $\eta(0) < 1$, then $\varphi^{<1} = \text{True}$. If $\eta(0) = 1$, then $\varphi^{<1} = \psi_2^{<1}$. Indeed, since η is strictly decreasing, the only chance of φ to have $\llbracket \pi, \varphi \rrbracket = 1$ is when both $\eta(0) = 1$ and $\llbracket \pi^0, \psi_2 \rrbracket = 1$. Since a satisfaction value cannot exceed 1, the latter happens iff $\eta(0) = 1$ and $\pi \not\models \psi_2^{<1}$ (where the “only if” direction is valid when π is a lasso, as is assumed).

Finally, it is easy to see that $|\varphi^+|$ and $|\varphi^{<1}|$ are both $O(|\varphi|)$.

A.2 The standard translation of LTL to AWA

For completeness, we bring here the construction of the translation from LTL to AWA, which we use in Theorem 1. For the correctness proof, see e.g. [27].

Given an LTL formula φ over the atomic propositions AP , we construct an AWA $\mathcal{A}_\varphi = \langle Q, 2^{AP}, Q_0, \delta, \alpha \rangle$ as follows. The state space Q consists of all the subformulas of φ , and their negations (we identify $\neg\neg\psi$ with ψ). The initial state is φ , and the accepting states are all the formulas of the form $\neg(\psi_1 U \psi_2)$. It remains to define the transition function.

We start with a few notations. For a Boolean formula θ over Q , we define its *dual formula* $\bar{\theta}$ by induction over the construction of θ , as follows.

- For $\psi \in Q$ we have $\bar{\bar{\psi}} = \neg\psi$.
- $\bar{\text{True}} = \text{False}$ and $\bar{\text{False}} = \text{True}$.
- $\bar{\psi_1 \vee \psi_2} = \bar{\psi_1} \wedge \bar{\psi_2}$ and $\bar{\psi_1 \wedge \psi_2} = \bar{\psi_1} \vee \bar{\psi_2}$

The transition function can now be defined as follows. Let $\psi \in Q$ and $\sigma \in 2^{AP}$.

- If $\psi = p \in AP$, then $\delta(\psi, \sigma) = \begin{cases} \text{True} & p \in \sigma, \\ \text{False} & p \notin \sigma, \end{cases}$
- If $\psi = \zeta_1 \vee \zeta_2$, then $\delta(\psi, \sigma) = \delta(\zeta_1, \sigma) \vee \delta(\zeta_2, \sigma)$,
- If $\psi = \neg\zeta$, then $\delta(\psi, \sigma) = \neg\delta(\zeta, \sigma)$,
- If $\psi = X\zeta$, then $\delta(\psi, \sigma) = \zeta$,
- If $\psi = \zeta_1 U \zeta_2$, then $\delta(\psi, \sigma) = \delta(\zeta_2, \sigma) \vee (\delta(\zeta_1, \sigma) \wedge \zeta_1 U \zeta_2)$.

A.3 Continuation of the Proof of Theorem 1

We continue the proof that is given in the main text, showing that the constructed AWA $\mathcal{A}_{\varphi, v}$ is indeed finite and correct.

We first make some notations and observations regarding the structure of $\mathcal{A}_{\varphi, v}$. For every state $(\psi > t)$ (resp. $(\psi < t)$) we refer to ψ and t as the state's *formula* and *threshold*, respectively. If the outermost operator in ψ is a discounting operator, then we refer to its discounting function as the state's *discounting function*. For states of Type-2 we refer to their *formula* only (as there is no threshold).

First observe that the only cycles in $\mathcal{A}_{\varphi, v}$ are self-loops. Indeed, consider a transition from state q to state $s \neq q$. Let ψ_q, ψ_s be the formulas of q and s , respectively. Going over the different transitions, one may see that either ψ_s is a strict subformula of ψ_q , or s is a Type-2 state, or both ψ_q and ψ_s have an outermost discounting operator with discounting functions η and η^{+1} respectively. By induction over the construction of φ , this observation proves that there are only self-cycles in $\mathcal{A}_{\varphi, v}$.

We now observe that in every run of $\mathcal{A}_{\varphi, v}$ on an infinite word w , every infinite branch (i.e., a branch that does not reach True) must eventually be in a state of the form $\psi_1 U \psi_2 > t$, $\psi_1 U \psi_2 < t$, $\psi_1 U \psi_2$ or $\neg(\psi_1 U \psi_2)$ (if it's a Type-2 state). Indeed, these states are the only states that have a self-loop, and the only cycles in the automaton are self-loops.

We start by proving that there are finitely many states in the construction. First, all the sub-automata that correspond to Type-2 states have $O(|\varphi|)$ states. This follows immediately from Lemma 1 and from the construction of an AWA from an LTL formula.

Next, observe that the number of possible state-formulas, up to differences in the discounting function, is $O(|\varphi|)$. Indeed, this is simply the standard closure of φ . It remains to prove that the number of possible thresholds and discounting functions is finite.

We start by claiming that for every threshold $t > 0$, there are only finitely many reachable states with threshold t . Indeed, for every discounting function $\eta \in \mathcal{D}$ (that appears in φ), let $i_{t, \eta} = \max \left\{ i : \frac{t}{\eta(i)} \leq 1 \right\}$. The value of $i_{t, \eta}$ is defined, since the functions tend to 0. Observe that in every transition from a state with threshold t , if the next state is also with threshold t , then the discounting function (if relevant) is either some $\eta' \in \mathcal{D}$, or η^{+1} . There are only finitely many functions of the former kind. As for the latter kind, after taking η^{+1} $i_{t, \eta}$ times, we have that $t/\eta^{+i_{t, \eta}}(0) > 1$. By the definition of δ , in this case the transitions are to the Boolean-formula states (i.e., True, False, or some ψ^+), from which there are finitely many reachable states. We conclude that for every threshold, there are only finitely many reachable states with this threshold.

Next, we claim that there are only finitely many reachable thresholds. This follows immediately from the claim above. We start from the state $\varphi > v$. From this state, there are only finitely many reachable discounting functions. The next threshold that can be encountered is either $1 - v$, or $\frac{v}{\eta(0)}$ for η that is either in \mathcal{D} or one of the η^{+i} for $i \leq i_{v, \eta}$. Thus, there are only finitely many such thresholds. Further observe that if a different threshold is encountered, then by the definition of δ , the state's formula is deeper in the generating tree of φ . Thus, there are only finitely many times that a threshold can change along a single path. So by induction over the depth of the generating tree, we can conclude that there are only finitely many reachable thresholds.

We conclude that the number of states of the automaton is finite.

Next, we prove the correctness of the construction. From Lemma 1 and the correctness of the standard translation of LTL to AWA, it remains to prove that for every path π and for every state $(\psi > v)$ (resp. $(\psi < v)$):

1. If $\llbracket \pi, \psi \rrbracket > v$ (resp. $\llbracket \pi, \psi \rrbracket < v$), then π is accepted from $(\psi > v)$ (resp. $(\psi < v)$).
2. If $\pi = u \cdot v^\omega$ and π is accepted from state $(\psi > v)$ (resp. $(\psi < v)$) then $\llbracket \pi, \psi \rrbracket > v$ (resp. $\llbracket \pi, \psi \rrbracket < v$).

The proof is by induction over the construction of φ , and is fairly trivial given the definition of δ .

A.4 Proof of Lemma 2

We start by defining *generalized Büchi automata*. An NGBA is $\mathcal{A} = \langle Q, \Sigma, \delta, Q_0, \alpha \rangle$, where Q, Σ, δ, Q_0 are as in NBA. The acceptance condition is $\alpha = \{F_1, \dots, F_k\}$ where $F_i \subseteq Q$ for every $1 \leq i \leq k$. A run r of \mathcal{A} is accepting if for every $1 \leq i \leq k$, r visits F_i infinitely often.

We now proceed with the proof.

Consider the AWA \mathcal{A} obtained from φ using the construction of Section 3.2.

In the translations of AWA to NBA using the method of [12], the AWA is translated to an NGBA whose states are the subset-construction of the AWA. This gives an exponential blowup in the size of the automaton. We claim that in our translation, we can, in a sense, avoid this blowup.

Intuitively, each state in the NGBA corresponds to a conjunction of states of the AWA. Consider such a conjunction of states of \mathcal{A} . If the conjunction contains two states $(\psi < t_1)$ and $(\psi < t_2)$, and we have that $t_1 < t_2$, then by the correctness proof of Theorem 1, it holds that a path π is accepted from both states, iff π is accepted from $(\psi < t_1)$. Thus, in every conjunction of states from \mathcal{A} , there is never a need to consider a formula with two different “<” thresholds. Dually, every formula can appear with at most one “>” threshold.

Next, consider conjunctions that contain states of the form $(\psi_1 U_\eta \psi_2 < t)$ and $(\psi_1 U_{\eta+k} \psi_2 < t)$. Again, since the former assertion implies the latter, there is never a need to consider two such formulas. Similar observations hold for the other discounting operators.

Thus, we can restrict the construction of the NGBA to states that are conjunctions of states from the AWA, such that no discounting operator appears with two different “offsets”.

Further observe that by the construction of the AWA, the threshold of a discounting formula does not change, with the transition to the same discounting formula, only the offset changes. That is, from the state $(\psi_1 U_\eta \psi_2 < t)$, every reachable state whose formula is $\psi_1 U_{\eta+k} \psi_2$ has threshold t as well. Accordingly, the possible number of thresholds that can appear with the formula $\psi_1 U_\eta \psi_2$ in the subset construction of \mathcal{A} , is the number of times that this formula appears as a subformula of φ , which is $O(|\varphi|)$.

We conclude that each state of the obtained NGBA is a function that assigns each subformula⁵ of φ two thresholds. The number of possible thresholds and offsets is linear in the number of states of \mathcal{A} , thus, the number of states of the NGBA is $|\mathcal{A}|^{O(|\varphi|)}$.

Finally, translating the NGBA to an NBA requires multiplying the size of the state space by $|\mathcal{A}|$, so the size of the obtained NBA is also $|\mathcal{A}|^{O(|\varphi|)}$.

A.5 Proof of Theorem 2

We construct an AWA $\mathcal{A} = \mathcal{A}_{\varphi, v}$ as per Section 3.2, with some changes.

Recall that the “interesting” states in \mathcal{A} are those of the form $\psi_1 U_{\eta+i} \psi_2$. Observe that for the function \exp_λ it holds that $\exp_\lambda^{+i} = \lambda^i \cdot \exp_\lambda$. Accordingly, we can replace a state of the form $\psi_1 U_{\exp_\lambda^{+i}} \psi_2 < t$ with the state $\psi_1 U_{\exp_\lambda} \psi_2 < \frac{t}{\lambda^i}$, as they express the same assertion. Finally, notice that $\exp_\lambda(0) = 1$. Thus, we can simplify the construction of \mathcal{A} with the following transitions:

Let $\sigma \in 2^{AP}$, then we have that

$$\begin{aligned}
 - \delta((\psi_1 U_{\exp_\lambda} \psi_2 > t), \sigma) &= \\
 &\begin{cases} \delta((\psi_2 > t), \sigma) \vee \\ [\delta((\psi_1 > t), \sigma) \wedge (\psi_1 U_{\exp_\lambda} \psi_2 > \frac{t}{\lambda})] & \text{if } 0 < t < 1, \\ \text{False} & \text{if } t \geq 1, \\ \delta(((\psi_1 U_{\exp_\lambda} \psi_2)^+), \sigma) & \text{if } t = 0. \end{cases} \\
 - \delta((\psi_1 U_{\exp_\lambda} \psi_2 < t), \sigma) &= \\
 &\begin{cases} \delta((\psi_2 < t), \sigma) \wedge \\ [\delta((\psi_1 < t), \sigma) \vee (\psi_1 U_{\exp_\lambda} \psi_2 < \frac{t}{\lambda})] & \text{if } 0 < t \leq 1, \\ \text{True} & \text{if } t > 1, \\ \text{False} & \text{if } t = 0. \end{cases}
 \end{aligned}$$

⁵ where a subformula may have several occurrences, e.g., in the formula $p \wedge Xp$ we have two occurrences of the subformula p

$$\begin{aligned}
- \delta((\psi_1 \mathbf{U}_{\text{exp}_\lambda} \psi_2 > t), \sigma) &= \\
&\begin{cases} \delta((\psi_2 > t), \sigma) \wedge \\ [\delta((\psi_1 > t), \sigma) \vee (\psi_1 \mathbf{U}_{\text{exp}_\lambda} \psi_2 > \frac{t}{\lambda})] & \text{if } 0 < t < 1, \\ \text{False} & \text{if } t \geq 1, \\ \delta(((\psi_1 \mathbf{U}_{\text{exp}_\lambda} \psi_2)^+), \sigma) & \text{if } t = 0. \end{cases} \\
- \delta((\psi_1 \mathbf{U}_{\text{exp}_\lambda} \psi_2 < t), \sigma) &= \\
&\begin{cases} \delta((\psi_2 < t), \sigma) \vee \\ [\delta((\psi_1 < t), \sigma) \wedge (\psi_1 \mathbf{U}_{\text{exp}_\lambda} \psi_2 < \frac{t}{\lambda})] & \text{if } 0 < t \leq 1, \\ \text{True} & \text{if } t > 1, \\ \text{False} & \text{if } t = 0. \end{cases}
\end{aligned}$$

The correctness and finiteness of the construction follows from Theorem 1, with the observation above. We now turn to analyze the number of states in \mathcal{A} .

For every state $(\psi > t)$ (resp. $(\psi < t)$) we refer to ψ and t as the state's *formula* and *threshold*, respectively. Observe that the number of possible state-formulas is $O(|\varphi|)$. Indeed, the formulas in the states are either in the closure of φ , or are of the form ψ^+ , where ψ is in the closure of φ . This is because in the new transitions we do not carry the offset, but rather change the threshold, so the state formula does not change.

It remains to bound the number of possible thresholds. Consider a state with threshold t . In every succeeding state⁶, the threshold (if exists) can either remain t , or change to $1 - t$ (in case of negation), or t/λ , where $\lambda \in F(\varphi)$, providing $t < 1$.

Initially, we ignore negations. In this case, the number of states that can be reached from a threshold t is bounded by the size of the set

$$\left\{ \prod_{i=1}^k \lambda_i : \prod_{i=1}^k \lambda_i > t, k \in \mathbb{N}, \forall i \lambda_i \in F(\varphi) \right\}$$

This length of the products can be bounded by $\log_\mu t = \frac{\log t}{\log \mu} = O(\log t)$, where $\mu = \max F(\varphi)$. Thus, the number of possible values is bounded by $(\log_\mu t)^{|F(\varphi)|}$. From here we denote $\log_\mu t$ by ℓ .

Next, we consider negations. Observe that in every path, there are at most $|\varphi|$ negations. In every negation, if the current state has threshold s , the threshold changes to $1 - s$. We already proved that from threshold t we can get at most $(\ell)^m$ states. Furthermore, every such state has a threshold of the form

$$\frac{t}{\prod_{i=1}^\ell \lambda_i}$$

where $\lambda_i \in F(\varphi) \cup \{1\}$ (we allow 1 instead of allowing shorter products).

Consider a negation state with threshold $s = \frac{t}{\prod_{i=1}^\ell \lambda_i}$. If $s = 1$ then in the next state the threshold is 0, and every reachable state is part of a Boolean AWA corresponding to an LTL formula, and thus has polynomially many reachable states.

Otherwise, we have that $s < 1$. Denote $F(\varphi) = \{\lambda_1, \dots, \lambda_m\}$, and assume that $t, \lambda_1, \dots, \lambda_m$ can be written as $t = \frac{p_t}{q_t}$ and for all i , $\lambda_i = \frac{p_i}{q_i}$ (which is possible as they are in \mathbb{Q}), then we have that

$\log(1 - s) = \log(1 - \frac{t}{\prod_{i=1}^\ell \lambda_i}) = \log(\prod_{i=1}^\ell \lambda_i - t) - \log(\prod_{i=1}^\ell \lambda_i) > \log(\prod_{i=1}^\ell \lambda_i - t)$, where the last transition is because $\log(\prod_{i=1}^\ell \lambda_i) < 0$.

Let $q_{\max} = \max \{q_1, \dots, q_m\}$, we now observe that

$$\begin{aligned}
\prod_{i=1}^\ell \lambda_i - t &= \prod_{i=1}^\ell \frac{p_i}{q_i} - \frac{p_t}{q_t} = \frac{\prod_{i=1}^\ell p_i q_t - \prod_{i=1}^\ell q_i p_t}{\prod_{i=1}^\ell q_i q_t} \geq \\
&\frac{1}{\prod_{i=1}^\ell q_i q_t} \geq \frac{1}{q_{\max}^\ell q_t}
\end{aligned}$$

⁶ This is almost correct. In fact, since δ is defined inductively, we may go through several transitions.

where the last transitions are because all the numbers are natural, and $s > 1$, so the numerator must be a positive integer. From this we get that

$$\log(1 - s) > \log\left(\frac{1}{q_{max}^\ell q_t}\right) = -(\ell \cdot (\log q_{max} + \log q_t))$$

We see that we can bound $1 - s$ from below by a rational number whose representation is linear in that of t and $|\langle \varphi \rangle|$. Denote this linear function by $f(n)$. Thus, continuing in this manner through $O(\varphi)$ negations, starting from the threshold v , we have that the number of thresholds reachable from every state is bounded by $\underbrace{(f(f(\dots f(\log v))))^m}_{O(\varphi)}$

which is single exponential in $|\langle \varphi \rangle|$ and the description of v .

We conclude that the number of states of the automaton is single-exponential.

A.6 Proof of Theorem 4

Recall that if π is a computation such that $\llbracket \pi, \varphi \rrbracket > v$, then $\mathcal{A}_{\varphi, v}$ accepts π . The converse however, is not true. Still, by carefully examining the construction in Theorem 3.2, we observe that if $\mathcal{A}_{\varphi, v}$ accepts a computation π , then $\llbracket \pi, \varphi \rrbracket \geq v$ (note the non-strict inequality).

Assume a partition of the letters in AP to input and output signals, denoted I and O , respectively. By following standard (Boolean) procedures for synthesis (see [24]), we can generate from $\mathcal{A}_{\varphi, v}$ a deterministic tree automaton \mathcal{D} that accepts a 2^O -labeled 2^I -tree iff all the paths along the tree are accepted in $\mathcal{A}_{\varphi, v}$. A tree that is accepted in this manner represents a transducer that realizes φ with value at least v . Accordingly, if there exists a transducer \mathcal{T} , all of whose computations satisfy $\llbracket \pi, \varphi \rrbracket > v$, then the tree that represents such a transducer is accepted in \mathcal{D} . Thus, the language of \mathcal{D} is non-empty, and a non-emptiness witness is a transducer (tree) \mathcal{T}' , all of whose computations satisfy $\llbracket \pi, \varphi \rrbracket \geq v$.

We remark that this solution is only partial, as there might be an I/O -transducer whose computations π all satisfy $\llbracket \pi, \varphi \rrbracket \geq v$, but we will not find it.

A.7 Translating $\text{PLTL}^{\text{disc}}[\mathcal{D}]$ formulas to 2AWA

As we now show, the construction we use when working with the “infinite future” U_η operator is similar to that of the one we use for the “finite past” S_η operator. The key for this somewhat surprising similarity is the fact our construction is based on a threshold. Under this threshold, we essentially bound the future that needs to be considered, thus the fact that it is technically infinite plays no role.

A 2AWA is a tuple $\mathcal{A} = \langle \Sigma, Q, q_0, \delta, \alpha \rangle$ where Σ, Q, q_0, α are as in AWA. The transition function is $\delta : Q \times \Sigma \rightarrow \mathcal{B}^+(\{-1, 1\} \times Q)$. That is, positive Boolean formulas over atoms of the form $\{-1, 1\} \times Q$, describing both the state to which the automaton moves and the direction in which the reading head proceeds.

As in $\text{LTL}^{\text{disc}}[\mathcal{D}]$, the construction extends the construction for the Boolean case. It is not hard to extend Lemma 1 and generate Boolean PLTL formulas for satisfaction values in $\{0, 1\}$.

Given a $\text{PLTL}^{\text{disc}}[\mathcal{D}]$ formula φ and a threshold $t \in [0, 1]$, we construct a 2AWA as in Theorem 1 with the following additional transitions:⁷

- $\delta((Y\psi > t), \sigma) = \langle -1, (\psi > t) \rangle$
- $\delta((Y\psi < t), \sigma) = \langle -1, (\psi < t) \rangle$.
- $\delta((\psi_1 S \psi_2 > t), \sigma) = \delta((\psi_2 > t), \sigma) \vee (\delta((\psi_1 > t), \sigma) \wedge \langle -1, (\psi_1 S \psi_2 > t) \rangle)$.
- $\delta((\psi_1 S \psi_2 < t), \sigma) = \delta((\psi_2 < t), \sigma) \wedge (\delta((\psi_1 < t), \sigma) \vee \langle -1, (\psi_1 S \psi_2 < t) \rangle)$.
- $\delta((\psi_1 S_\eta \psi_2 > t), \sigma) =$

$\delta((\psi_2 > \frac{t}{\eta(0)}), \sigma) \vee [\delta((\psi_1 > \frac{t}{\eta(0)}), \sigma) \wedge \langle -1, (\psi_1 S_{\eta+1} \psi_2 > t) \rangle]$	if $0 < \frac{t}{\eta(0)} < 1$,
False	if $\frac{t}{\eta(0)} \geq 1$,
$\delta(((\psi_1 S_\eta \psi_2)^+), \sigma)$	if $\frac{t}{\eta(0)} = 0$.

⁷ In addition, the atoms in the transitions in Theorem 1 are adjusted to the 2AWA syntax by replacing each atom q by the atom $\langle 1, q \rangle$.

$$\begin{aligned}
& - \delta((\psi_1 \mathbf{S}_\eta \psi_2 < t), \sigma) = \\
& \begin{cases} \delta((\psi_2 < \frac{t}{\eta(0)}), \sigma) \wedge [\delta((\psi_1 < \frac{t}{\eta(0)}), \sigma) \vee \langle -1, (\psi_1 \mathbf{S}_{\eta+1} \psi_2 < t) \rangle] & \text{if } 0 < \frac{t}{\eta(0)} \leq 1, \\ \text{True} & \text{if } \frac{t}{\eta(0)} > 1, \\ \text{False} & \text{if } \frac{t}{\eta(0)} = 0. \end{cases}
\end{aligned}$$

The correctness of the construction and the analysis of the blowup are similar to those in Section 3.

A.8 Correctness Proof of the Construction in Section 5.3

Consider the case of $(\psi_1 \mathbf{O}_{\eta, z} \psi_2 > t)$ (the dual case is similar). We check when this assertion holds.

First, if $z = t$, then this assertion is equivalent to $\psi_1 \mathbf{U} \psi_2 > t$ (this follows directly from the semantics). Thus, assume that $z \neq t$.

If $\tau < 0$, then $t < (1 - \eta(0))z$, so in particular, for every value of $\llbracket \pi^0, \psi_2 \rrbracket$, it holds that $t < \eta(0)\llbracket \pi^0, \psi_2 \rrbracket + (1 - \eta(0))z$, so the assertion is true, since the first operand in the sup is greater than t .

If $\tau \geq 1$ then $\eta(0) + (1 - \eta(0))z \leq t$, so both $\eta(0)\llbracket \pi^0, \psi_2 \rrbracket + (1 - \eta(0))z \leq t$ and $\eta(0)\llbracket \pi^0, \psi_1 \rrbracket + (1 - \eta(0))z \leq t$. Thus, every operand in the sup has an element less than t , so the sup cannot be greater than t , so the assertion is false.

If $0 \leq \tau < 1$, then similarly to the case of \mathbf{U}_η - the assertion is equivalent to the following: either $\psi_2 > \tau$, or both $\psi_1 > \tau$ and $(\psi_1 \mathbf{O}_{\eta+1, z} \psi_2 > t)$.

It remains to show that there are still only finitely many states. Since $z \neq t$, we get that $\lim_{\eta(0) \rightarrow 0} \tau = \begin{cases} \infty & t > z \\ -\infty & t < z \end{cases}$.

Thus, after a certain number of transitions, τ takes a value that is not in $[0, 1]$, in which case the next state is True or False, so the number of states reachable from $\varphi > t$ is finite.

We remark that this is not true if $z = t$, which is why we needed to treat this case separately.